HUMAN PERFORMANCE ON NEGATIVE-SLOPE SCHEDULES OF REINFORCEMENT: A TEST OF REINFORCEMENT MAXIMIZATION

By

ERIC A. JACOBS

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

1997

ACKNOWLEDGMENTS

I wish to thank the many individuals who helped me complete this dissertation. Firstly, I wish to recognize my family for all the moral and financial support they have given me over these many years. I am especially indebted to my wife, Angela Rabb Jacobs, for sharing the joys and sacrifices inherent in the pursuit of an academic career. Her unwavering confidence in my abilities has maintained my own. I am grateful to Tim Hackenberg for all the guidance and assistance he has given me. I also wish to acknowledge the other members of my supervisory committee: Marc Branch, Henry Pennypacker, Don Stehouwer, Brian Iwata, and Mark Lewis. I have learned much from the instruction and challenge of these learned scholars. I also wish to thank my fellow students for the many hours of intellectual melee. I particularly wish to acknowledge Manish Vaidya, Cindy Pietras, Dave Stafford, Diana Walker, Karen Anderson and Kevin Jackson for all the conversations, favors, and good times. I am also appreciative of Warren K. Bickel for giving me the opportunity to continue in academics and for his patience during the compilation of this manuscript. Lastly, I wish to express my undying appreciation of E. F. Malagodi for introducing me to the worldview of radical behaviorism. Although I have learned many things from Ed, of most importance has been the value of a good teacher.

TABLE OF CONTENTS

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

HUMAN PERFORMANCE ON
NEGATIVE-SLOPE SCHEDULES OF REINFORCEMENT:
A TEST OF REINFORCEMENT MAXIMIZATION

By

ERIC A. JACOBS

August, 1997

Chairman: Timothy D. Hackenberg, Ph.D.
Major Department: Psychology

Previous research has demonstrated that the behavior of normal adult humans is more sensitive to temporally remote outcomes than that of nonhumans under laboratory conditions. It has been suggested these between-species differences may be the result of human verbal behavior. The present work attempted to minimize verbal influences by using procedures that are sufficiently subtle to preclude accurate description. In Experiment 1, panel pressing was maintained in five adults by schedules of points exchangeable for money. Following exposure to variable-interval 30-second and linear-variable-interval 30-second schedules, the subjects were exposed to various negative-slope schedules that establish an inverse relationship between response rate and reinforcement rate. Specifically, responses produced points according to an interval schedule and

canceled point deliveries according to a ratio schedule. The way to earn the most points was to respond at a rate of approximately 2 per minute. Overall response rates varied as a function of negative-slope schedule parameters, but overall reinforcement rate was not maximized in any of the negative-slope conditions. Bimodal distributions of interresponse times preceding point delivery and similar within-session patterns of responding were obtained for four of the subjects. In Experiment 2, sensitivity to the inverse relationship between response rate and reinforcement rate was assessed by first exposing two adults to negative-slope schedules, then to yoked-linear-variable-interval schedules that produced the same temporal distribution of point deliveries. Overall reinforcement rate was again not maximized in the negative-slope conditions. Response rates were approximately the same or lower in the yoked-linear-variable-interval condition, suggesting that behavior was insensitive to the contingencies of the negative-slope conditions. The results, however, are difficult to interpret due to order effects and to the absence of response patterning characteristic of steady-state performance on negative-slope schedules. In postexperimental questionnaires, none of the subjects in either experiment were able to quantify the relationship between responding and point delivery. Overall, the results show that when verbal intrusions are minimized, human behavior may be more susceptible to immediate consequences, and more consistent with nonhuman behavior, than has been demonstrated in prior research.

INTRODUCTION

The degree to which behavior is sensitive to short-term versus longer-term consequences is an area of research with important implications for basic and applied researchers. For the basic researcher, such studies can provide critical information about the time scale over which consequences are functionally related to behavior, which is important in understanding how reinforcement creates and maintains behavioral units. For the applied researcher, such studies may be useful in understanding behavioral problems that stem from excessive control by proximal consequences. For example, people abuse drugs, gamble excessively, incur large amounts of debt, and eat poorly without seeming regard for the temporally deferred consequences of these activities. The ability to demonstrate and manipulate similar "impulsive" behavior patterns in a controlled laboratory setting may greatly increase our understanding of the conditions under which they occur, which may, in turn, make it possible to correct or prevent the occurrence of such behavior outside of the laboratory.

The most widely used procedures to assess control by immediate versus long-term consequences involve mutually exclusive choices between a small immediate reinforcer and a larger delayed reinforcer (Ainslie, 1974; Logue, 1988; Rachlin & Green, 1972). When behavior is governed by the smaller, immediate reinforcer, the performance is said to be "impulsive" or to demonstrate "impulsivity." Conversely, when behavior is

governed by the larger delayed reinforcer, the performance is said to be "self-controlled" or to demonstrate "self-control."

One of the advantages of such procedures is that they have been used profitably with human and nonhuman subjects alike. In a comparative review of the "self-control" literature, Logue (1988), however, described some general disparities between the performances of humans and nonhumans on procedures involving recurring choices between an immediate smaller reinforcer and a temporally deferred larger reinforcer. With few exceptions, nonhumans consistently choose the smaller immediate reinforcer, thereby minimizing the delay to reinforcement, whereas verbal adult humans consistently choose the larger delayed reinforcer, thereby maximizing overall reinforcement density. Thus, human performances are said to be "self-controlled," and nonhuman performances are said to be "impulsive." On the face of it, such results suggest that human behavior may be more sensitive than nonhuman behavior to temporally deferred outcomes.

The tendency for human behavior to be more sensitive than nonhuman behavior to long-term outcomes is not limited to the area of self-control, however. Procedures examining choice in situations of diminishing returns have also revealed quantitatively discernible differences between species in sensitivity to temporally remote consequences. Originally introduced by Hodos and Trumbule (1967) with chimpanzees, the procedure involves repeated choices between a fixed schedule, the requirements of which remain constant within a session, and a progressive schedule, the requirements of which are initially lower than the fixed but gradually become more stringent with each reinforcement delivered by that schedule.

The dependent measure of most interest under these procedures is the pattern of switching from the progressive schedule to the fixed. Under some conditions, short-term consequences (those based solely upon minimizing the delay to the upcoming reinforcer) and long-term consequences (those based upon maximizing overall reinforcement rate) support different choice patterns. The short-term consequences support consistent switching from the progressive schedule to the fixed when the schedule values are equal. Overall reinforcement frequency is maximized, however, by selecting the fixed schedule earlier in the progression, when its schedule requirements greatly exceed those of the progressive schedule.

Subjects consistently switch from the progressive to the fixed schedule in advance of the point at which the schedule values are equal, indicating that behavior is sensitive to consequences beyond the current choice cycle. This effect has been reported with rhesus monkeys, (Hineline & Sodetz, 1987), pigeons, (Hackenberg & Hineline, 1992; Mazur & Vaughan, 1987; Wanchisen, Tatham, & Hineline, 1988) and humans (Hackenberg & Axtell, 1993; Jacobs & Hackenberg, 1996; Wanchisen, Tatham, & Hineline, 1992), in addition to the original study with chimpanzees (Hodos & Trumbule, 1967). As with self-control procedures, however, human performances on these procedures are generally better described by accounts based upon the maximization of reinforcement density, whereas the performances of nonhumans are generally better characterized by a delay-based account (Hackenberg & Axtell, 1993; Mazur & Vaughan, 1987). Although the results of these choice experiments have revealed that nonhuman behavior can come under control of temporally removed consequences to some extent, they also suggest that

human behavior may be more sensitive than nonhuman behavior to long-term, overall outcomes.

The between-species disparity in the results of the "self-control" experiments described above may reflect a fundamental difference between humans and other species in the sensitivity to deferred consequences. It is difficult, however, to establish unequivocally that differences in performance are solely of phylogenic origin. To do so, one would have to rule out at least two other possibilities: (a) that species differences were not due to differences in procedure; or (b) that species differences were not due to differences in preexperimental histories. These will be considered in turn.

Of the many procedural differences between experiments with humans and those with nonhumans, differences in the types of reinforcers may contribute to between-species performance differences. Studies with nonhumans typically use unconditioned reinforcers (e.g., grain presentation to food deprived pigeons), whereas experiments with humans typically use token reinforcers (points exchangeable for money). In "self-control" experiments, for example, humans make repeated choices between a small amount of points exchangeable for money delivered immediately and a larger amount of points delivered after a brief delay. The delay to exchange the points for money is usually fixed and remains the same (i.e., at session's end) regardless of the choice patterns. Humans may consistently prefer the deferred larger amount of points because there is no immediate advantage to obtaining points quickly, as they are not exchanged for consumable reinforcers until later. This hypothesis is supported by the fact that impulsive patterns in humans are more likely with reinforcers that are more immediately

effective, such as escape from white noise (Navarick, 1982; Solnick, Kannenberg, Eckerman, & Waller, 1980) or access to edible reinforcers (Logue & King, 1991). Moreover, recent experiments by Hyten, Madden, and Field (1994) and by Jackson and Hackenberg (1996) have demonstrated that similar performances in people and pigeons, respectively, can be achieved by manipulating exchange delays in "self-control" procedures involving token reinforcers.

Even when procedural differences are minimized, the possibility remains that human operant performance in the laboratory may interact with a subject's prior verbal history (Catania, Mathews, & Shimoff, 1982; Horne & Lowe, 1993; Rosenfarb, Newland, Brannon, & Howey, 1992). Humans not only respond to contingencies, but describe those contingencies and how they are related to behavior. Skinner (1969) introduced a distinction between rule-governed and contingency-shaped behavior to help characterize the different ways in which human behavior comes under control of the environment. Contingency-shaped behavior is behavior that is governed by direct experience with contingencies. Rule-governed behavior, on the other hand, is behavior that is under the control of complex antecedent stimuli established through a history of socially mediated reinforcement (Cerutti, 1989; Horne & Lowe, 1996; Skinner, 1957). Normal adult humans enter the laboratory with the ability not only to follow rules, but also to derive rules describing the relationship between behavior and consequences. Verbal humans may augment the effects of contingencies with verbal descriptions of those contingencies; as a result, their performances may differ from those of nonverbal subjects.

Like nonverbal behavior, rule-following and rule-generating are classes of behavior established and maintained by consequences. For example, such abilities may be helpful for establishing responding in another member of the verbal community, in occasioning a complicated performance that may have been otherwise forgotten, or in generalizing behavior to contingencies that bear some similarity to those previously experienced. Skinner (1969) suggested that another possible benefit of rule-governed behavior is to bring behavior into contact with temporally deferred outcomes. According to his view, rules may be helpful in sustaining performances that would normally be curtailed by more proximal consequences. This may account for the between-species differences in performance observed in experiments that place long-term and short-term consequences in opposition.

There is some evidence in the literature to support this position. Adult humans in a "self-control" experiment, for example, reported using maximization strategies that were derived by counting during the intervals (Logue, Peña-Correal, Rodriguez, & Kabela, 1986). That is, the subjects were able to describe the overall relationship between their choices and the amount of reinforcement. Such descriptions may have functioned as verbal discriminative stimuli that brought their behavior into contact with the deferred consequences.

Additional evidence of the role of verbal behavior in establishing control by deferred consequences comes from a study of age-related differences in self-control by Sonuga-Barke, Lea, and Webley (1989). In that study, girls ages 4, 6, 9, and 12 made repeated choices between a small reinforcer (one token) and a larger reinforcer (two tokens). The

delay to the small reinforcer was held constant at 10 s, whereas the delay to the large reinforcer was varied from 20 s to 50 s, across conditions. Overall reinforcement maximization was thus contingent upon preferring the smaller, more immediate, reinforcer when the delay to the larger was greater than 20 s. The results indicated developmental differences that corresponded with the subjects' degree of verbal sophistication. The behavior of the four-year-olds was most similar to that of nonhumans. They tended to be either indifferent or to prefer the shorter delay to the smaller reinforcer across all conditions. The six-and nine-year-olds tended to prefer the larger, delayed reinforcer across all conditions, even when doing so decreased the overall amount of tokens obtained. The twelve-year-olds showed increasing preference for the smaller reinforcer as the delay to the large increased, thus approximately maximizing the overall amount of tokens earned. Postsession verbal reports corresponded with nonverbal performances. The six-and nine-years-olds reported preferring the larger, delayed reinforcer regardless of the schedule values, whereas the twelve-year-olds reported preferring whichever alternative would produce the most tokens overall. Although it is difficult to establish the precise role of verbal behavior in generating these performances, the overall pattern of results suggests that verbal behavior may interact with sensitivity to long term outcomes.

There are limits, however, on the extent to which a description can bring behavior into contact with contingencies. Although someone may derive a rule that brings behavior into contact with a given contingency of reinforcement, the rule-governed performance is unlikely to be thoroughly equivalent to a contingency-shaped performance (Skinner,

1969). Training someone to drive a car with a manual transmission provides an example of the use of rule governance to bring behavior under the control of a particular set of contingencies. The novice driver is often given instructions on how to operate the clutch and the accelerator simultaneously. The instructor may say, "Slowly release the clutch while simultaneously depressing the accelerator." Although such instruction is likely to be of some assistance in generating behavior sufficient to move the car, it is highly unlikely that a well executed performance will occur on the first attempt. A fluid departure will occur only after sufficient experience with the contingencies imposed by the car itself. The instructor cannot offer a precise description of appropriate performance based upon his or her everyday interaction with the device. "Slowly" is a rather vague description that is unlikely to occasion smooth, well-executed operation of the clutch on the first attempt. On the other hand, the listener is unlikely to respond appropriately to even the most accurate of instructions under such circumstances. The instructor could take the time and effort to measure precisely the relationship between the operation of the clutch and the accelerator and to refine the instruction given to the pupil to "Lift your left foot off the clutch at a rate of 2 cm per second, while depressing the accelerator with your right foot at a rate of 1 cm per sec." Without extensive training, however, it is unlikely that the pupil will be able to follow this instruction with any greater degree of success than the original instruction.

It is likely that similar limits of description and interpretation occur when the speaker and listener occupy the same skin. That is, there may be limits to the influence of self-instruction as well. Expert athletes and artists often cannot describe their own

performances in great detail. Their behavior is shaped and established through an extensive history of contact with the contingencies of the respective activities. Consider the act of driving a golf ball. The distance and the direction the ball moves when hit is determined by subtle aspects of the operation of the club--so subtle that they may elude precise description; one may be left with only vague metaphorical extensions that are of little help in the instruction of others, or oneself.

It is not altogether surprising that human subjects in laboratory experiments can often provide accurate descriptions of the programmed contingencies. Take, for example, the laboratory self-control experiment. Subjects are repeatedly exposed to simple dichotomous choices between alternatives that have rather discriminable outcomes after relatively short delays. These are precisely the conditions under which one might expect accurate verbal descriptions to be brought to bear on nonverbal responding.

Although it is probably not possible to generate a human performance under standard laboratory conditions that does not occasion some verbal behavior with respect to the contingencies, it may be possible to minimize the role of verbal mediation by using contingencies in which the relationship between responding and its consequences is very subtle and difficult to describe accurately. To that end, the present study was undertaken to assess humans' sensitivity to long-term consequences using procedures unlikely to give rise to accurate descriptions. The procedures were patterned after an experiment conducted by Vaughan and Miller (1984), designed to test interpretations of performances based upon maximization of reinforcement density. In that study, pigeons' keypecking was maintained on schedules of reinforcement that established an inverse

relationship between overall response rate and overall reinforcement rate. On this type of schedule, reinforcement density is maximized by responding relatively slowly.

Under these procedures, reinforcers are scheduled according to a linear variable interval (LVI) schedule. LVI schedules differ from traditional variable interval (VI) schedules in that upon completion of an interval a reinforcer is added to a "store" of available reinforcers and timing of the ensuing interval begins immediately. Reinforcement follows every response when the "store" is positive, and 1 is subtracted from the "store" following each reinforcement. So, for example, if there were 3 reinforcers in the "store," the next 3 responses would be reinforced and the store would be depleted, thus requiring another interval to expire before reinforcement was again available. One way to portray the overall differences between VI and LVI schedules is through the use of molar feedback functions. Molar feedback functions are quantitative descriptions of the relationship between overall rate of responding and overall rate of reinforcement established by a schedule. Figure 1 shows the feedback functions for both an LVI 30-s schedule (solid line) and for a traditional VI 30-s schedule (dashed line). Whereas the VI function asymptotically approaches the programmed reinforcement rate across the entire range of response rates, the LVI function is constant across all but the lowest response rates. LVI schedules are thus more forgiving of deviations in response rate than traditional VI schedules, in that reinforcement density is maximized for all response rates greater than or equal to the inverse of the schedule value (i.e., 2.0 responses per min in Figure 1).

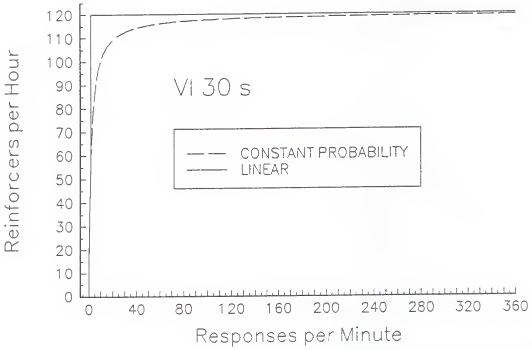

Figure 1. Molar feedback functions for linear and constant-probability VI 30 -s schedules (See text for details).

An inverse relationship between overall response rate and overall reinforcement rate can be established by conjointly subtracting reinforcers from the "store" according to a fixed-ratio (FR) schedule. The "store" can be driven into a negative state following periods of relatively high rate responding, thus requiring periods of relatively low rate responding before the value of the "store" is again positive and reinforcement can occur. Figure 2 shows the resulting feedback functions when reinforcers are conjointly scheduled by an LVI 30-s schedule and canceled every 100, 50, or 10 responses (from top to bottom, respectively). The functions are negatively sloped across all response rates that are greater than the inverse of the LVI value. For all subtraction ratios, overall reinforcement density is maximized at the point on the function where response rate equals 1/VI. Beyond this point, faster rates of responding result in decreasing rates of reinforcement until the rate of reinforcement is effectively 0.0 per min. The rate at which the reinforcement rate decreases as a function of response rate is inversely related to the FR value. Note the differences in the slopes of the functions in Figure 2. Reinforcement rate decreases sharply over a narrower range of response rates when reinforcement is canceled every 10^{th} response than when it is canceled every 50^{th} or every 100^{th} response.

The feedback function provides a continuum along which performance can be evaluated with respect to overall reinforcement density. The overall response rate that maximizes overall reinforcement rate will always be the inverse of the LVI schedule value (See Figure 2). Because any significant deviation from this rate will result in a

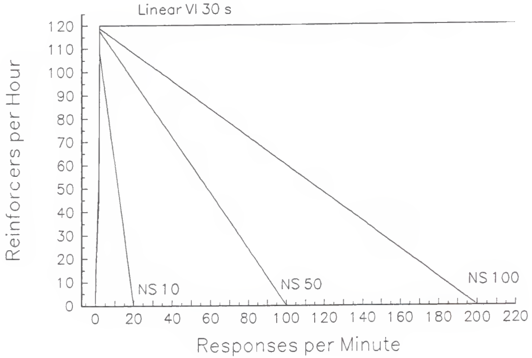

Figure 2. Molar feedback functions for LVI 30-s and negative slope schedules with subtraction ratios of 10, 50, and 100.

corresponding decrease in overall reinforcement density, deviations from reinforcement maximization are readily detectable.

To the extent that maximization of overall reinforcement density takes place at the expense of increases in local delays to reinforcement, the negative slope procedures share some common ground with “self-control” procedures. As response rate decreases on negative slope schedules and approaches the inverse of the LVI schedule value, a greater percentage of the intervals scheduling reinforcement will be shorter than the time between successive responses, or interresponse times (IRTs). As IRT length increases (i.e., responses rate decreases), the probability that the store is positive also increases. Any relative approximation to the rate at which reinforcement density is maximized will therefore incur longer local delays to reinforcement. Thus, as with “self-control” procedures, there is competing control on negative slope schedules by a high probability of immediate reinforcement and increases in overall reinforcement density.

Unlike "self-control" procedures, however, the aperiodicity of point delivery established by the schedule may decrease the likelihood that subjects will accurately describe the contingency. Although subjects may be able to report that going "slow" produces more points overall than going "fast," it is unlikely that they will be able to offer a precise quantitative description of the relationship between overall response rate and overall rate of point delivery. If verbal behavior enhances sensitivity to remote consequences, as some research appears to suggest, then perhaps limiting accurate verbal descriptions of the contingencies will also limit nonverbal sensitivity to those contingencies. That is, we may expect to see more "suboptimal" performance (deviations

from maximization) than has been reported in prior studies of human behavior under laboratory conditions.

The present investigation consisted of two experiments. In Experiment 1, panel pressing was maintained in five adult human subjects by schedules of points exchangeable for money. Following exposure to traditional VI 30-s and LVI 30-s schedules, the subjects were exposed to various negative-slope schedules of reinforcement to assess the degree to which their behavior was sensitive to these contingencies. In Experiment 2, sensitivity to the presence of the subtraction contingency was assessed by first exposing subjects to negative slope schedules and then to linear-variable-interval schedules with the same temporal distribution of point deliveries. Performances from the two conditions were then compared to determine if responding was sensitive to the presence of the subtraction contingency. It was hoped that the inverse relationship between responding and point delivery established in the negative slope conditions was sufficiently subtle to preclude accurate description, thereby minimizing verbal mediation.

The present study was also concerned with providing a more detailed picture of response patterning on negative slope schedules. To that end, daily cumulative response records and IRTs preceding point delivery were collected and analyzed in relation to the contingencies. Thus, apart from the theoretical issues that prompted the use of negative slope schedules to begin with, the study was also aimed at a more thorough understanding of this interesting but seldom used contingency arrangement.

EXPERIMENT 1

Method

Subjects

Two male and three female adult humans participated fully in Experiment 1 in exchange for money. The subjects were recruited via a classified advertisement in a campus newspaper. None of the subjects had previously or was currently enrolled in coursework in behavior analysis or learning theory. The subjects were informed prior to their first session that they would be paid $1.50 per session and that they could receive a bonus by earning points worth $0.05 each. In order to encourage full participation, collection of bonus earnings was contingent upon completion of the study. The subjects were also instructed not to bring personal items (e.g., food, tools, smoking materials, portable radios, or time pieces) into the work space and that violation of this rule would result in dismissal from the study. Overall earnings (including bonuses) ranged from $3.00 to $7.16 per hour (median = $6.00 per hour).

Six additional subjects were recruited for the study but did not participate beyond the initial condition. Five of these subjects withdrew from the study and another was dismissed due to insufficient baseline response rates. The data from these subjects will not be considered here.

Apparatus

The manipulandum consisted of the right side of a 12.5 cm X 7.5 cm dual response panel, illuminated with red light and mounted centrally on a BRS/Foringer human operant panel (HTC-603). A computer monitor was seated on top of the panel and a keyboard was located in front of the panel. Data were recorded and contingencies controlled by a program compiled in Quickbasic® running on a Compuadd® model 286 computer. Cumulative response records were collected using a Gerbrands® cumulative recorder.

Procedure

The following instructions were read to the subject prior to the first session and were displayed on the computer monitor prior to the start of each session.

> Please read carefully. You have already earned $1.50. In order to receive bonus earnings, you will have to earn points by pressing the red panel. Each point you earn is worth $0.05. So, for example, if you earn 80 points then your bonus would be $4.00. You may leave the room at any time in the event of an emergency. Thank you for your participation.

Sessions were divided into 10-min blocks and began when the subject typed his or her subject number on the keyboard and pressed the enter key. Each block was separated by a rest period, during which time the manipulandum was darkened and no programmed contingencies were in effect. Rest periods were terminated by a response on the computer keyboard. Sessions were comprised of five such blocks for Subject 111, and 3

blocks for the other four subjects. Sessions were generally conducted on weekdays at approximately the same time of day.

Experimental conditions

Table 1 contains the sequence of conditions and the number of sessions conducted under each for the five subjects who completed Experiment 1. Conditions were changed when mean response rates and within-session patterns of responding were deemed stable via visual inspection of graphical representations of the data. The following are descriptions of each of the 4 types of experimental conditions.

Variable interval (VI) 30 s. Panel pressing was maintained on a variable-interval (VI) 30-s schedule of point delivery. The VI schedule consisted of 50 intervals generated using the method described by Fleshler and Hoffman (1962). The intervals were randomized and were then presented sequentially both within and across sessions. Upon completion of an interval, the next response would result in point delivery. Point delivery was signaled by a brief tone and by incrementing a counter that was continuously displayed on the computer monitor throughout the block. The next interval in the sequence began timing immediately after point delivery.

Linear variable interval (LVI) 30-s. Conditions were the same as in the VI 30 -s condition except that, upon completion of an interval, a point was "stored" in a bank of available reinforcers and the next interval began timing immediately. If a press occurred and the value of the store was positive, one point was delivered and one was subtracted from the store. In the absence of responding, the number of available reinforcers in the

Table 1. The sequence of conditions (and the number of sessions conducted under each) for all subjects under variable interval (VI), linear variable interval (LVI), and negative slope (NS) schedules of point delivery.

Subject				
S111	S888	S999	S211	S521
VI (36)	VI (8)	VI (10)	VI (5)	VI (11)
LVI (11)	LVI (5)	LVI (13)	CONJT (8)	LVI (7)
NS 326 (35)	NS 203 (6)	NS 54 (8)	VI (5)	NS 221 (8)
NS 41 (10)	NS 136 (5)	NS 30 (9)	LVI (5)	NS 183 (5)
NS 20 (13)	NS 100 (6)	NS 22 (10)	NS 253 (7)	NS 156 (8)
	NS 88 (7)		NS 138 (6)	NS 139 (17)
	NS 74 (5)		NS 105 (28)	NS 117 (25)
	NS 58 (6)		NS 33 (29)	
	NS 47 (7)			
	NS 10 (9)			

store would continue to escalate with the completion of each successive interval. Subjects were not informed of the value of the store, which carried across blocks and sessions, nor of the change in conditions.

Negative slope (NS) FR n. In this condition, points were added to the store according to an LVI 30 s schedule as described above, but every *n*th response subtracted one point from the store. If points were subtracted more frequently than they were set up by the LVI schedule, the store could be driven into a negative state. Both the value of the store and the position within the FR were carried across sessions. The subtraction ratios for individual subjects are listed in Table 1. These values were based upon the mean response rates from the last five sessions in the preceding condition, and were selected such that the overall rate of point delivery would be reduced by one half if the steady-state response rates of the preceding condition prevailed. In the first negative-slope condition for Subject 111, for example, the subtraction ratio value was 326, equal to the mean responses per min from the last 5 sessions of the LVI 30 s condition. If this subject continued to respond at 326 responses per min in the NS 326 condition, then one point per minute would be subtracted from the store according to the FR while the LVI 30 s schedule was simultaneously adding 2 points per min to the store. The resulting net reinforcement rate would be approximately one point per minute, or approximately half the steady-state rate of point delivery from the LVI condition.

The subtraction ratios of Subjects 888 and 521 were determined slightly differently. The initial subtraction ratio was selected such that the overall rate of point delivery would be 0.5 per minute instead of 1.0 per minute. Both of these subjects showed little

sensitivity to the negative slope schedules after several conditions. To produce decreases in response rates in the NS 58 and NS 47 conditions for Subject 888 and in the NS 139 and NS 117 conditions for Subject 521, subtraction ratios that reduced the rate of point delivery to approximately 0.0 per minute were used. In the final negative slope (NS 10) condition for Subject 888, the targeted rate of point delivery was -6.3 per minute.

Conjoint VI 30 s (IRT < 2 s):Random ratio (IRT < 2 s). (CONJT VI:RR) For Subject 211, the rate of responding following the initial 5 session exposure to the VI condition was insufficient to progress to the remaining conditions. To rectify this, the subject was exposed to a complex schedule of point delivery in which interresponse times (IRTs) less than 2 s were reinforced according to random ratio (RR) and VI 30-s schedules of point delivery, conjointly. On average, points were delivered following IRTs less than 2 s every 30 s by the VI schedule and after every n^{th} such IRT by the RR schedule, where n equals the random ratio schedule value. The third 10-min block of the session served as a probe during which the random ratio schedule was suspended and points were scheduled solely by the VI 30-s schedule with the IRT requirement. The random-ratio schedule value was initially set at 30, but was decreased across the first 3 sessions of the condition until response rates increased, indicating schedule control by a response-based contingency. The ratio requirement was then increased gradually over 4 sessions. This manipulation established high rates of responding when the ratio requirement was relatively low and eventually sustained those rates with reinforcement densities comparable to that produced by efficient VI 30-s performance when the ratio requirement was relatively large. At that point, this schedule was discontinued and the VI condition

was reimplemented. To maintain overall session earnings comparable to the other conditions, sessions were comprised of 3 blocks that terminated after 10-min or 20 point deliveries, whichever came first.

Procedural irregularities

There were a few deviations from the procedure during the course of the experiment. Subject 211 did not respond during the first session. Prior to the second session, the experimenter reread the instructions to the subject and gestured toward the appropriate response panel. For Subjects 111, 888, and 999, the final condition was concluded arbitrarily. For Subjects 111 and 999, the final condition ended due to schedule conflicts. Subject 888 left town during a semester break and did not return to the experiment.

Results

Molar Feedback Functions

Figure 3 shows the mean rate of point delivery versus the mean response rate from the last 5 sessions under all conditions for each subject. These points have been superimposed upon the feedback functions for the respective schedules. The uppermost curve in each graph is the molar feedback function for the LVI condition and the dashed curve below it is the molar feedback function for the VI condition. The diamonds and the triangles in Figure 3 show the data obtained under the LVI and the VI conditions, respectively. Rate of point delivery increased following the transition to the LVI from the VI conditions for all subjects, as indicated in the figure by a slight vertical displacement of the diamonds relative to the positions of the triangles; this result

Figure 3. Mean rate of point delivery versus the mean response rates from the last 5 sessions of the VI 30-s, LVI 30-s, and NS conditions for each subject. The data are superimposed on the molar feedback functions for each condition. Numbers near the abscissae indicate the subtraction ratios of the NS conditions. Note individually scaled axes.

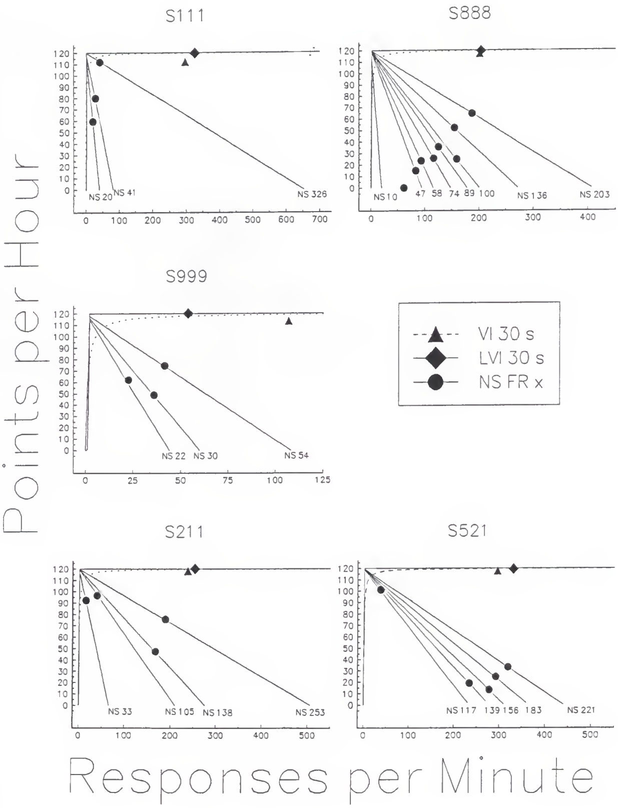
S111
S888
S999
S211
S521
VI 30 s
LVI 30 s
NS FR x
NS 20
NS 41
NS 326
NS 10
47
58
74
89
100
NS 136
NS 203
NS 22
NS 30
NS 54
NS 33
NS 105
NS 138
NS 253
NS 117
139
156
183
NS 221
Points per Hour
Responses per Minute

is in accord with the programmed changes in the contingencies across the two conditions, and occurred whether response rates increased (Subjects 111, 211, & 521), decreased (Subject 999), or remained the same (Subject 888).

The remaining curves in the graphs of Figure 3 are the molar feedback functions for the negative-slope conditions. The functions are negatively sloped across all response rates greater than 2.0 per minute. Beyond this point, faster rates of responding result in decreasing rates of reinforcement until the rate of point delivery reaches a minimum of 0.0 per min. The rate at which rate of point delivery decreases as a function of response rate is inversely related to the FR value. The filled circles in Figure 3 indicate the obtained rate of point delivery versus the obtained response rate under the negative slope conditions. If reinforcement density was consistently maximized, the points would converge at the apexes of the feedback functions, that is, at 2 responses per min. The response rates of all the subjects across all negative-slope conditions exceeded this value and, therefore, reinforcement rate deviated from the maximum obtainable.

Sensitivity to changes in the contingencies can be assessed by comparing the obtained steady-state values to those encountered earlier in the condition. An appreciation of the situation confronting the subjects following a transition from one condition to the next can be achieved by drawing a vertical line from one of the obtained points in Figure 3 to the abscissa. The point at which this line intersects the feedback function of the subsequent condition will indicate the rate of reinforcement encountered immediately following the transition, assuming similar response rates are obtained. Any steady-state data that are plotted up and to the left of this intersection on the feedback function

indicate a reduction in response rate and a corresponding increase in reinforcement rate. For example, the largest absolute change in response rate occurred with Subject 111 during exposure to the first negative-slope condition (NS 326). A subtraction ratio of 326 was selected such that an overall reinforcement rate of 60 points per hour would be obtained if there were no changes in overall response rate. The rate decreased across the NS 326 condition from 326 responses per minute at the conclusion of the LVI condition to 41 responses per minute at the conclusion of the negative-slope condition. The decrease in response rate resulted in a steady-state reinforcement rate that was comparable to that obtained in the VI condition. The obtained point is plotted up and to the left of the targeted intersection (e.g., 326 responses per minute, 60 points per hour) and, thus, indicates sensitivity to differences in the contingencies across the two conditions.

In Figure 3, complete insensitivity to changes in the FR requirement for reinforcement omission is evidenced by a vertical arrangement of the points from the two conditions. In other words, the response rate would remain the same, as the reinforcement rate changed across conditions according to the different FR values. Across subjects and negative-slope conditions, the only instance of complete insensitivity to changes in the FR subtraction requirement, occurred under NS 100 for Subject 888. The overall rate of responding under NS 100 was similar to that obtained under NS 136. This resulted in a considerable decrease in overall reinforcement rate across the two conditions.

Overall response rate decreased relative to that of the previous condition in the remaining 22 negative-slope conditions across subjects. Although these decreases

indicate sensitivity to the contingencies, the changes in response rate seldom produced increases in the steady-state rates of point delivery across conditions. The steady-state rates of point delivery were lower than those of the previous conditions in 18 of 23 negative-slope conditions across subjects. The five exceptions to this were Conditions NS 89 for Subject 888, NS 22 for Subject 999, NS 105 for Subject 211, and NS 139 and NS 117 for Subject 521. Steady-state rate of point delivery increased across consecutive negative slope conditions only once across subjects (Conditions NS 117 & NS 139 for Subject 521).

The point obtained for Subject 888 under NS 10 deserves comment because it does not fall on the positive region of the molar feedback function for that condition. This subject was exposed to these contingencies for nine sessions and, despite receiving only one point in the first session of the condition, the response rate remained sufficiently high that the rate of point omission exceeded the rate of point availability. By the end of these nine sessions the value of the point "store" was -1424. For the value of the "store" to again become positive, this subject would have had to cease responding entirely for 24 consecutive sessions. The subject left town at this point for a semester break and never returned to the experiment.

Session-by-Session Patterns

Figures 4, 5, & 6 show the mean response rates (filled circles) and mean rates of point delivery (open circles) across sessions for all conditions. Whereas only steady-state data were portrayed in Figure 3, these figures show transition data and trends in response and reinforcement rates.

Figure 4. The mean response rates (filled circles) and mean rates of point delivery (open circles) across sessions for all conditions for Subjects 111 (top panel) and 888 (bottom panel). Note individually scaled axes.

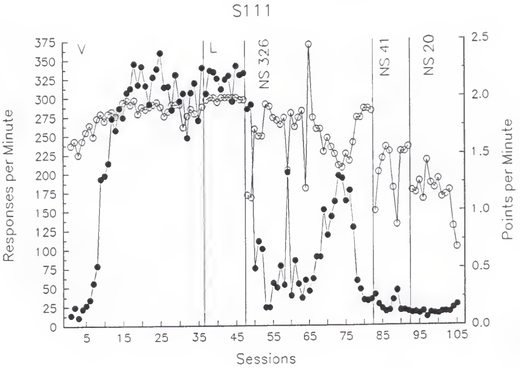
S111
V
L
NS 326
NS 41
NS 20
Responses per Minute
Points per Minute
Sessions
375
350
325
300
275
250
225
200
175
150
125
100
75
50
25
0
2.5
2.0
1.5
1.0
0.5
0.0
5 15 25 35 45 55 65 75 85 95 105

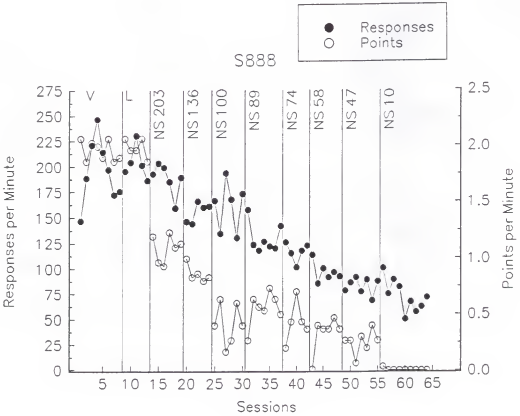
Responses
Points
S888
V
L
NS 203
NS 136
NS 100
NS 89
NS 74
NS 58
NS 47
NS 10
Responses per Minute
Points per Minute
Sessions
275
250
225
200
175
150
125
100
75
50
25
0
2.5
2.0
1.5
1.0
0.5
0.0
5 10 15 20 25 30 35 40 45 50 55 60 65

For Subject 111, (top panel, Fig. 4) the rate of responding in the VI condition was initially low, but increased across sessions before reaching a steady state of approximately 300 responses per minute. Rate of point delivery varied directly with response rate throughout the VI condition, and asymptotically approached 2.0 points per minute as the response rate increased. Rates of responding and point delivery were comparable to the steady-state rates obtained under the VI condition throughout the LVI condition, although there was somewhat less variation.

Response rates decreased relative to those from the VI and LVI conditions under the negative-slope conditions. At the conclusion of the NS 326 condition, response rates reached a steady state of approximately 41 responses per minute and the resulting rates of point delivery were comparable to those obtained under the VI condition. Under the NS 41 condition, the response rates again decreased, but the decreases were slight in comparison to those of the preceding condition. As a result, the rates of point delivery were lower than those obtained at the conclusion of the NS 326 condition. In the final NS condition for this subject (NS 20), rates of responding were similar to those of the steady-state rates of the preceding condition in all but the final 2 sessions, during which response rates increased slightly. As a result, the rates of point delivery were lower overall than the steady-state rates obtained under the NS 41 condition throughout the condition, and decreased considerably in the last two sessions. Unfortunately, this subject could no longer participate due to a scheduling conflict, and data collection for this subject concluded with this slight increase in response rate.

The bottom panel of Figure 4 shows the mean response rates and the mean rates of point delivery across sessions for all conditions for Subject 888. Under the VI condition, response rates increased across the first 4 sessions to a condition maximum, then decreased to a steady state of approximately 180 responses per minute. Although there was some between-session variability in the overall rates of point delivery, there were no trends across the condition as a whole. Under the LVI condition, response rates increased slightly relative to the steady state rates of the VI condition but the rates of point delivery were similar to those obtained in the VI condition.

With the exception of NS 100 condition, response rates decreased slightly with each decrease in the subtraction ratio in the negative-slope conditions. Under the NS 100 condition, the overall rates of responding were comparable to that of the preceding condition, but there was an increase in the variability. The overall rates of point delivery also decreased across the NS conditions in 7 of 8 cases. The exception to this pattern occurred in the NS 89 condition, as described earlier.

For Subject 999 (top panel, Fig. 5), response rates increased across the first 5 sessions of the VI condition before decreasing to a steady state of approximately 95 responses per minute. The rates of point delivery varied around a mean of 1.89 points per minute throughout the condition. Response rates decreased throughout the LVI condition and reached a steady state of approximately 54 responses per minute by its conclusion. The rates of point delivery were higher under the LVI condition relative to those of the VI condition, and there was greater between session variability in the rates as well.

Figure 5. The mean response rates (filled circles) and mean rates of point delivery (open circles) across sessions for all conditions for Subjects 999 (top panel) and 211 (bottom panel). Note individually scaled axes.

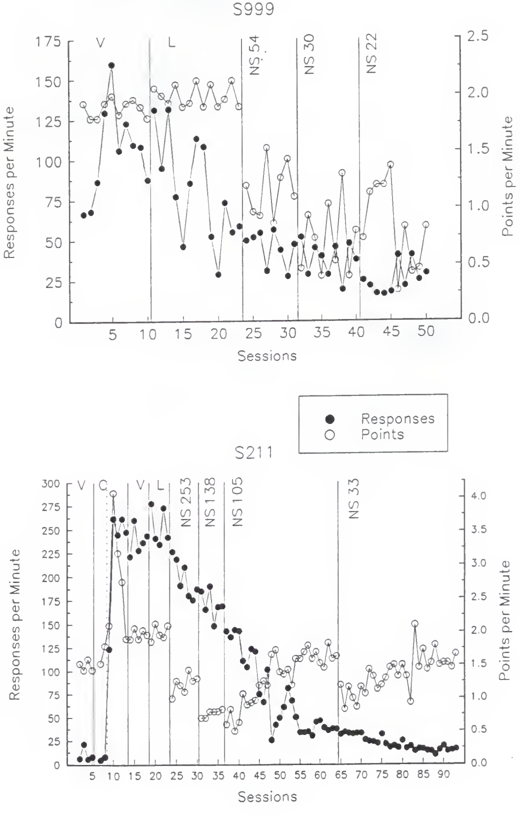
S999
V
L
NS 54
NS 30
NS 22
Responses per Minute
Points per Minute
Sessions
Responses
Points
S211
V
C
V
L
NS 253
NS 138
NS 105
NS 33
Responses per Minute
Points per Minute
Sessions

Under the negative-slope conditions, response rates and rates of point delivery decreased with each successive decrease in the subtraction ratio. Under the NS 54 condition, the response rates decreased slightly relative to the rates of the preceding LVI condition and, as a result, there was a large decrease in rate of point delivery. In addition, between-session variability in rates of point delivery increased under the NS 54 condition. In the subsequent condition (NS 30), there was a slight decrease in response rates relative to the previous condition, although there was considerable overlap in the range of response rates from the two conditions. The rates of point delivery in the NS 30 condition were lower than the rates from the previous condition because the change in response rates across conditions was so slight. Response rates initially decreased under the NS 22 condition and, as a result, the rates of point delivery initially increased. As the condition progressed, however, response rates increased slightly, producing decreases in the rates of point delivery by the conclusion of the condition. Unfortunately, this subject, too, had to leave the experiment due to a scheduling conflict.

The bottom panel of Figure 5 shows mean response rates and mean rates of point delivery across sessions for Subject 211. Due to low response rates in the initial 5 sessions of the VI condition, this subject was exposed to the Conjoint schedule described above. The data from Session 6 were lost due to a programming error. Upon return to the VI condition, response rates decreased slightly, but were sufficient to continue the experiment. Response rates increased under the LVI condition, but the rates of point delivery in the VI and LVI conditions were similar.

This subject was exposed to four negative slope conditions, during which response rates decreased within and across conditions. Rates of point delivery decreased following the transition from one negative slope condition to the next, but increased within each condition. The largest absolute change in response rates occurred during the third negative slope condition (NS 105); as a result, the steady state rate of point delivery exceeded those of the previous two conditions (NS 138 & NS 253). Response rates decreased further in the final negative slope condition (NS 33) and, following an initial decrease, the rates of point delivery increased to a steady state that was similar to that of the NS 105 condition.

Figure 6 shows the mean response rates and the mean rates of point delivery across sessions for Subject 521. Under the VI condition, response rates increased across sessions, reaching a steady state of approximately 305 responses per minute. With the exception of the first session, the rate of point delivery varied around a mean of 1.80. Response rates increased in the LVI condition, reaching a steady state of approximately 333 responses per minute. The rate of point delivery under the LVI condition was similar to that of the VI condition, but there was somewhat more variability.

Subject 521 was exposed to 5 negative slope conditions. Across the first 3 of these, response rates decreased slightly with each decrease in the subtraction ratio. Rates of point delivery in the first negative-slope condition decreased considerably relative to those obtained in the LVI, then decreased slightly across the two subsequent conditions.

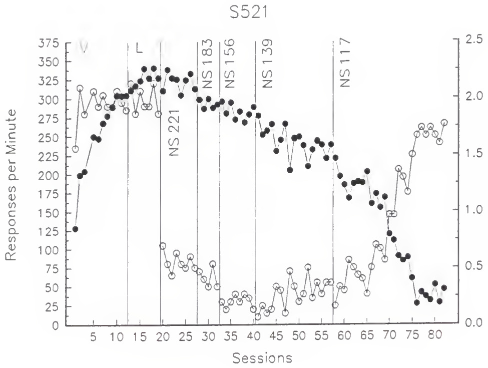

Figure 6. The mean response rates (filled circles) and mean rates of point delivery (open circles) across sessions for all conditions for Subject 521.

The largest absolute decrease in response rate occurred within the final condition (NS 117). The rate of point delivery increased considerably within the condition to reach a level nearly comparable to that of the baseline VI conditions.

Within-Session Patterns

With the exception of Subject 999, the within-session response rates were similar across subjects. Portions of representative cumulative records from sessions conducted under VI, LVI, and NS conditions are shown in Figure 7 for Subject 211, selected to represent the other subjects. Each panel is the segment of the cumulative record from the second 10 min block of the session. The leftmost segment (Panel 1) is from Session 5 during the first exposure to the VI condition. The slope of the line is shallow and roughly constant across the 10 min block, indicating a low but steady rate of responding. The segment in Panel 2 is from Session 17 during the second exposure to the VI condition (following the Conjoint condition). Note that the slope of the line in this segment is steeper than that of the curve on the left, indicating a higher rate of responding. The next segment is from Session 23 during the LVI condition. The slope of this curve is similar to that from the preceding VI condition, indicating that there was little change in the response rates across the two conditions.

Terminal performance under the negative slope schedules was characterized by alternating periods of high-rate responding and pausing. The development of this response pattern can be seen in the remaining panels, taken from sessions conducted under negative slope conditions. In Panel 4 (NS 253, Session 30), occasional pauses

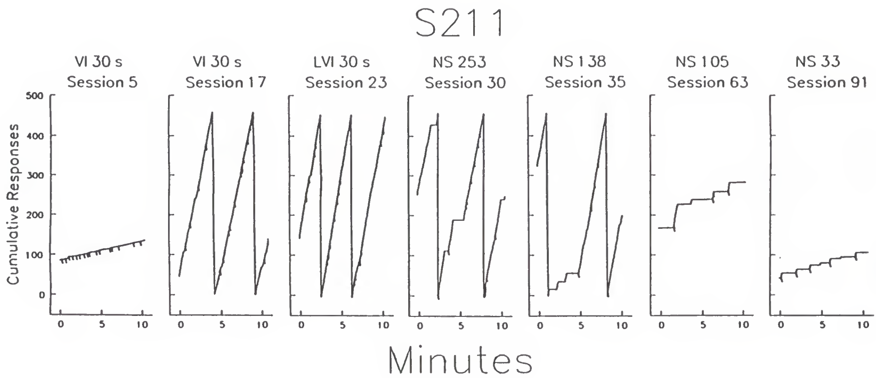

Figure 7. Portions of representative cumulative response records from sessions conducted under VI, LVI, and NS conditions for Subject 211. Each panel is the segment of the cumulative record from the second 10 min block of the session.

interrupt periods of high-rate responding. In the subsequent condition (NS 138, Session 35), the high-rate patterns were similar, but pausing was more frequent. Panel 6 (NS 105, Session 63) shows response patterning that is characteristic of terminal negative slope schedule performance: pauses alternate with periods of high-rate responding. The high-rate response "runs" frequently began with a series of point deliveries as the "store" was depleted and typically continued for many responses thereafter. Occasionally, point delivery occurred in the latter portion of these response runs. This "pause-run" pattern is even more pronounced and consistent in Panel 7 (NS 33, Session 91). Note that the "pauses" are more frequent and the "runs" are shorter.

The within-session response patterns of Subject 999 differed from those of the other subjects. In the VI condition, the overall rate of responding was lower than those of the other subjects and there was an aperiodic alternation between accelerating and decelerating response rates, producing a "wavy" cumulative response curve. These general characteristics persisted throughout the experiment. In the LVI condition, the overall rates decreased and momentary fluctuations in response rate became even more pronounced. Under successive negative slope conditions, the overall rate of responding decreased and pauses became more frequent, but the momentary changes in response rate described above persisted. Only in the final negative slope condition was there any evidence of the possible development of a "pause and run" pattern of responding, but even here, this pattern was less pronounced than for the other subjects. Instead, sporadic pausing was often preceded by decelerating response rates and followed by accelerating response rates.

Figure 8. Scatter plots of IRTs preceding point delivery across sessions for Subjects 211 (top panel) and 888 (bottom panel). Note individually scaled and logarithmic axes.

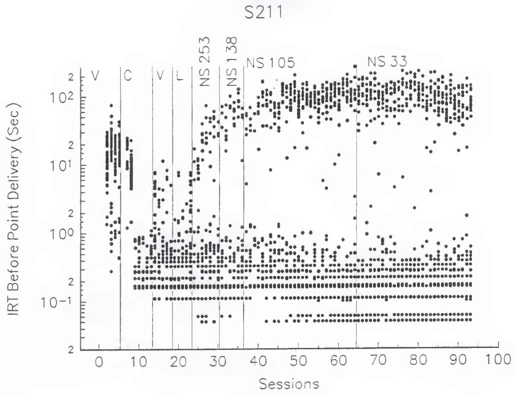
S211
V
C
V
L
NS 253
NS 138
NS 105
NS 33
IRT Before Point Delivery (Sec)
Sessions

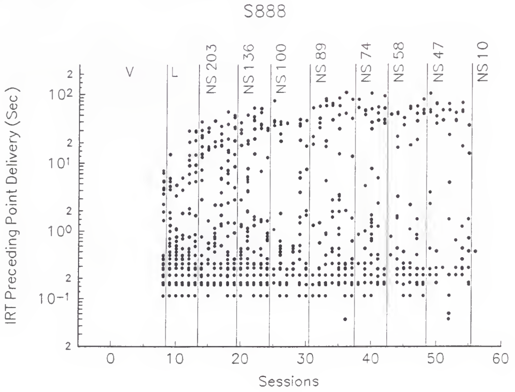
S888
V
L
NS 203
NS 136
NS 100
NS 89
NS 74
NS 58
NS 47
NS 10
IRT Preceding Point Delivery (Sec)
Sessions

Figure 9. Scatter plots of IRTs preceding point delivery across sessions for Subjects 521 (top panel) and 999 (bottom panel). Note individually scaled and logarithmic axes.

S521

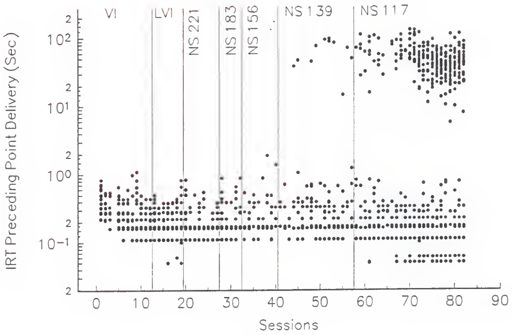
VI
LVI
NS 221
NS 183
NS 156
NS 139
NS 117
IRT Preceding Point Delivery (Sec)
10²
10¹
10⁰
10⁻¹
0
10
20
30
40
50
60
70
80
90
Sessions

S999

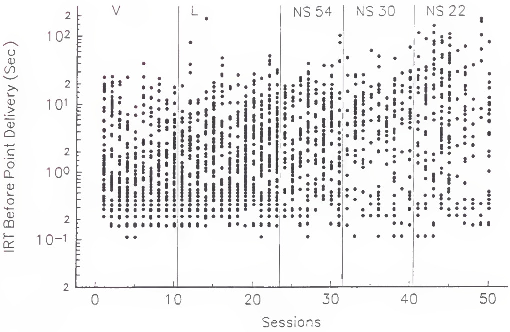
V
L
NS 54
NS 30
NS 22
IRT Before Point Delivery (Sec)
10²
10¹
10⁰
10⁻¹
0
10
20
30
40
50
Sessions

IRTs Preceding Point Delivery

The "pause-run" pattern of responding under negative slope conditions apparent in the cumulative records suggested that the IRT preceding reinforcement may be a useful measure to track across conditions. Figures 8 and 9 are scatter plots of the interresponse times preceding point delivery across sessions for all subjects except 111. (Subject 111, the first subject run in the experiment, was the subject whose data suggested the present measure.) Consistent with the "pause-run" patterns of the cumulative records, by the conclusion of the negative-slope conditions, the distributions of IRTs preceding point delivery were bimodal for three of these four subjects (888, 211, & 521). For Subject 999, although there was an increase in the relative frequency of longer IRTs across the negative slope conditions, the development of a bimodal distribution was not as pronounced.

Discussion

The general pattern of results suggests that the behavior of all five subjects was sensitive to the contingencies of the negative slope conditions. This was evidenced by changes in the steady-state rates of responding and point delivery, in the within-session response patterns depicted in cumulative response records, and in the distributions of IRTs preceding point delivery. Steady-state rates of responding decreased relative to those of the preceding condition with each successive decrease in the subtraction ratio in 22 of 23 cases across subjects (See Figure 3). Rates of point delivery, however, were not maximized in any of the negative slope conditions, suggesting that sensitivity to the inverse relationship between response rate and reinforcement rate was limited. The rates

of point delivery tended to increase as response rates decreased within negative slope conditions. Overall, however, the steady-state rates of point delivery generally decreased across successive negative slope conditions. The steady-state rates of point delivery exceeded those of the previous condition in only 5 of 23 negative slope conditions across subjects.

Changes in the within-session patterns of responding also suggest sensitivity to the changes in the reinforcement contingencies. For Subjects 111, 888, 211, and 521, constant, moderate rates of responding established in the VI and LVI conditions gave way to a pattern in which periods of pausing alternated with periods of high-rate responding under the negative slope conditions. For Subject 911, the within-session response patterns were characterized by low rates of responding that fluctuated unsystematically from moment to moment. The overall rates for this subject were, however, lower in the negative slope conditions than in the LVI or VI conditions.

The changes in the temporal distribution of responses depicted in the cumulative records were also reflected in changes in the distributions of IRTs preceding point delivery. The distributions of three of the four subjects (211, 888, & 521) for which this measure was recorded were bimodal, with one cluster of short IRTs (e.g., less than 1 s) and another cluster of longer IRTs (e.g., greater than 45 s). For Subject 911, the distributions of IRTs preceding point delivery were not bimodal, but the frequency of relatively long IRTs increased under the negative slope conditions.

Some compelling evidence of a limitation in sensitivity to the inverse relationship between overall rate of point delivery and overall response rate comes from sessions in

which unusually long pauses were followed by substantial increases in reinforcement rate. On four occasions across subjects, the end of a 10-min block was preceded by a particularly lengthy pause (e.g., in excess of 4 min). This occurred twice for Subject 999 (Sessions 36 & 38) and once each for Subjects 111 and 211 (Sessions 64 and 82, respectively). All subjects were under negative slope conditions at the time. For Subjects 111 and 211, these pauses preceded the conclusion of the final block of the session. Both of these subjects reported having fallen asleep during this period. While they were asleep, the program continued to "store" points as the intervals expired and the "stored" points were carried over to the following session. The subsequent sessions thus began with a series of FR 1 point deliveries as the "store" was depleted. As a result, the overall rates of point delivery for those sessions were 2.46 and 2.10 per minute for Subject 111 and 211, respectively (See Figures 4 & 5). For Subject 999, one pause preceded the conclusion of the first 10 min block of Session 36 and the other occurred prior to the conclusion of the second block of Session 38. The rest periods following these pauses were also unusually long, suggesting that he, too, had fallen asleep. As with the other subjects, the computer stored points during these pauses and the points were delivered successively at the start of the next block. As a result, the overall rates of point delivery during these sessions were the highest for the condition (See Figure 5). Interestingly, these experiences did not appear to influence immediately the overall rate or pattern of responding in the subsequent session or block.

Sensitivity to the contingencies of the negative slope schedule was apparent earlier in the sequence of conditions for Subject 111 than for the other subjects. The response rates

of Subject 111 decreased substantially during Session 50, the third session of the first negative slope condition (See Figure 4). The "pause and run" response pattern was also observed for the first time in the cumulative record from this session. The other subjects required much longer exposure to the negative slope contingencies before rate decreases of such magnitude were observed, if they occurred at all. These between-subject differences may have been due to a procedural difference that lessened the within-session correlation between rate of responding and point delivery for these subjects. For Subjects 888, 999, 211, and 521 sessions consisted of 3 10-min blocks, whereas, for Subject 111, sessions consisted of 5 such blocks. Subject 111 experienced each of the 50 intervals comprising the VI 30 s schedule about twice per session. Due to the decrease in session time, the other subjects experienced most intervals once and only repeated 10 intervals per session on average. Therefore, although the session-wide mean interpoint interval varied around 30 s across sessions, the mean for any given session may not have equaled 30 s for these subjects. This resulted in greater variability in overall rate of point delivery under the VI and LVI conditions and lessened the within-session correlation between response rate and point rate under the NS conditions, particularly at relatively large subtraction ratios. The influence of this weaker correlation is apparent in Figures 4, 5, and 6. One might expect that the inverse relationship between overall response rate and overall rate of point delivery would be readily apparent throughout the negative slope conditions in these figures. That is, relatively high rates of responding should produce relatively low rates of point delivery, and vice-versa. Only the data from Subject 111,

however, show this pattern robustly. The relationship is not as pronounced in the data from the negative-slope conditions for the remaining subjects.

Like the humans in the present study, the performance of nonhumans on negative slope schedules is not well characterized by reinforcement density maximization. Figure 10 (reprinted from Vaughan, 1982) shows the steady-state response rates from the negative slope conditions in Vaughan and Miller's (1984) experiment superimposed upon the molar feedback functions for the respective schedules. The response rates differed substantially from the inverse of the LVI values, and the reinforcement rates frequently approached 0 per minute. Despite substantial experience with these contingencies, low response rates were not obtained and overall reinforcement density was not maximized. These result have been subsequently replicated with rats (Ettinger, Reid, & Staddon, 1987; Reed & Schachtman, 1989, 1991).

Comparing the steady-state data presented in Figure 3 to those presented in Figure 10 from the Vaughan and Miller (1984) study suggests that the performances of humans are somewhat similar to those of pigeons. In the Vaughan and Miller (1984) study, each of three groups of pigeons was exposed to a particular subtraction ratio at LVI schedule values of 30, 45, and 90 s. The relevant data for comparison purposes, therefore, are the filled circles plotted on the upper most curves in the panels of Figure 10. These are the data from conditions in which the LVI schedule value was 30 s and the subtraction ratios were 20, 40, and 60, from left to right, respectively. Although interpretation is limited by differences in experimental design (Vaughan and Miller used a between-groups design,

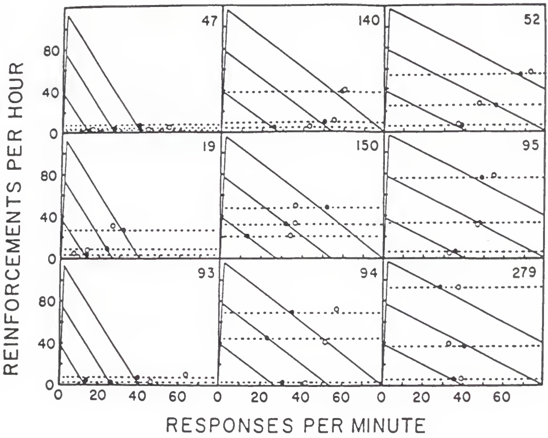

Figure 10. Reinforcement rates versus response rates of pigeons under NS (filled circles and solid lines) and yoked LVI (open circles and dashed lines) conditions from Vaughan and Miller (1984). The data have been plotted on the molar feedback functions for the respective conditions. From "Choice and the Rescorla-Wagner model." by W. Vaughan, Jr., 1982, in M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), Quantitative analyses of behavior: Vol. 2. Matching and maximizing accounts, (pp. 263 -279), Cambridge, MA: Ballinger. Copyright 1982 by Lawrence Erlbaum Associates, Inc. Reprinted by permission.

the present study a within-subject design), the data in Figure 10 suggest that rates of responding and reinforcement were inversely related to the subtraction ratio value. This was frequently observed with the human subjects as well (See Figure 3).

There were two instances in the present study, however, that were not consistent with this pattern. In the NS 105 condition for S211 and the NS 117 condition for S521, response rates decreased sufficiently to produce substantial increases in the rates of point delivery relative to those of conditions with larger subtraction ratios (See Figure 3). These deviations may have been driven by shifts in the within-session patterns of responding. The large absolute changes in rates of point delivery within these negative slope conditions were correlated with the development of the "pause and run" pattern of responding. Similarly, although it was the first negative slope condition for Subject 111, large absolute changes in rates of point delivery occurred in the NS 326 condition, where the "pause and run" pattern first developed.

Although cumulative response records were not presented for the pigeons in the Vaughan and Miller (1984) study, it is less likely that such patterns would have develop because a minimum 5 s interreinforcement interval was imposed. That is, if there was more than one reinforcer "stored", pecks within 5 s after reinforcer presentation would not produce another reinforcer. This contingency was omitted from the present study to enhance the inverse relationship between response rates and reinforcement rates, thereby providing a more stringent test of reinforcement density maximization.

EXPERIMENT 2

That human subjects consistently failed to maximize reinforcement density under the negative slope schedules in Experiment 1 raises the question of whether or not their behavior was sensitive to the inverse relationship between response and reinforcement rates. Instead, their performances may have been governed solely by the obtained overall rate of reinforcement. That is, it is possible that the pattern of reinforcement generated under the negative slope schedules influenced the overall rate and pattern of responding in a manner similar to that of a variable interval schedule with the same temporal distribution of reinforcement; the subtraction ratio may have exerted only an indirect influence by determining that distribution.

Herrnstein (1970) introduced a hyperbolic functional relation to account for steady-state response rates on VI schedules. It is based upon the premise of the matching law which states that the relative rate of responding equals the relative rate of reinforcement produced by that behavior. The expression takes the general form of

$$\underline{B} = \underline{k}(\underline{r}/(\underline{r} + \underline{r}_o)) \quad (1)$$

where $\underline{B}$ equals response rate, $\underline{r}$ equals the obtained rate of reinforcement produced by $\underline{B}$, $\underline{k}$ equals the maximum value of $\underline{B}$, and $\underline{r}_o$ equals the reinforcement rate corresponding to 1/2 $\underline{B}$ (Bradshaw, et al., 1976); $\underline{r}_o$ is theoretically equal to the rate of unscheduled reinforcement produced by behavior other than $\underline{B}$. With the aid of the two free

parameters, Herrnstein (1970) was able to fit this curve to data obtained by Catania and Reynolds (1968) with pigeons responding on a variety of VI schedules.

The data from the negative slope conditions in the Vaughan and Miller (1984) experiment suggest that overall response rate might match overall reinforcement rate. The filled circles in Figure 10 show the obtained response rates versus the obtained reinforcement rates from negative slope conditions. Although none of the pigeons was exposed to enough conditions to discern a hyperbolic relationship, response rates varied directly with reinforcement rates for 6 of the 9 birds, an outcome that is consistent with accounts based upon relative reinforcement rate (Herrnstein, 1970; Prelec, 1982). This outcome suggests that the pigeons' behavior might not have been sensitive to the subtraction ratio of the negative slope schedule.

The overall pattern of results from Experiment 1, however, indicates that human performance on similar negative-slope schedules was at least partially sensitive to the subtraction contingency. With the possible exception of Subject 888, the distributions of steady-state data in Figure 3 are not well characterized by hyperbolic functions. In the coordinates of Figure 3, the function described above would be plotted as a hyperbola that is concave upward, intersecting the ordinate at 0 responses per minute. The function would asymptotically approach some maximum rate of responding along the abscissa. According to the equation above, relatively higher rates of reinforcement should result in relatively higher rates of responding. Although this relationship holds for 18 of 23 negative-slope conditions across subjects, fairly pronounced decreases in response rate were accompanied by increases in reinforcement rate in four of the five exceptions. The

data from these conditions are not in accord with a straightforward application of Equation 1.

Experiment 2 was conducted to assess explicitly sensitivity to the relationship between rates of responding and reinforcement under negative slope schedules with human subjects. One way to assess sensitivity to such a relationship is to compare performance under these conditions to performance under a variable-interval schedule with the same temporal distribution of reinforcement. For example, subjects might first be exposed to a negative-slope schedule, then to a linear variable-interval schedule in which the reinforcement rate is yoked to that obtained in the preceding negative-slope condition. If the performances in the two conditions are similar, then it suggests that the overall rate of reinforcement is the relevant controlling factor and that responding is insensitive to the subtraction contingency. If response rates are consistently lower under the negative slope schedule than under the variable interval schedule, then it implies that the subjects' behavior is indeed sensitive to the inverse relationship between responding and reinforcement.

The results of such a manipulation can be seen in Figure 10. In the Vaughan and Miller (1984) study, the negative slope conditions alternated with LVI conditions in which the overall reinforcement rates equaled those of the preceding negative-slope conditions. The data from the LVI conditions (open circles) have been superimposed upon the molar feedback functions for those schedules (dashed lines) in the figure. Although not a particularly robust effect, the response rates tended to be higher in the

LVI conditions than in the associated negative slope conditions, indicating some degree of sensitivity to the negative slope contingencies.

The tendency for higher response rates under the yoked LVI schedules than under negative slope schedules with comparable reinforcement rates has been replicated in a series of experiments by Reed and Schachtman (1989, 1991) with rats as experimental subjects. In two of these experiments between-subject designs were used. Animals were paired according to similarities in their baseline performances. One rat in each pair was then arbitrarily assigned the master subject, exposed to a negative slope schedule of food delivery, and the other was assigned the yoked subject, exposed to an LVI schedule in which the reinforcement rates were yoked to those experienced by the master subject. In one study, reinforcement was yoked solely in terms of the overall rate, but in the other, the temporal distribution of the reinforcement was matched across the two subjects (Reed & Schachtman, 1989 & 1991, respectively). In the latter study, the roles of the subjects were later reversed, such that each subject had experience as both a master subject and a yoked subject. In both experiments, rats in the yoked groups tended to respond at slightly higher rates than the master rats that experienced the negative slope schedules; however, between subject differences in performance made interpretation of the results tenuous, at best.

In another experiment by Reed and Schachtman (1991), the comparison between performances under negative-slope and yoked LVI schedules was accomplished in the context of a within-subject design. Food was delivered according to a negative slope schedule in one component of a multiple schedule of reinforcement and according to a

yoked LVI in another. The component schedules were presented in strict alternation following each food delivery. Response rates were higher in the yoked components than in the negative slope components for all subjects, suggesting that their behavior was indeed sensitive to the inverse relationship between response rate and reinforcement rate in the negative slope components.

Experiment 2 extended the use of the yoked-control procedure to human subjects to assess sensitivity to the subtraction contingency. Using a reversal design, each subject was exposed to negative-slope and yoked-LVI schedules of points exchangeable for money. The interpoint intervals obtained under the negative slope condition were later presented in the same sequence in the yoked LVI condition. VI 30 s and LVI 30 s conditions served as the baseline conditions to which subjects were exposed prior to both the negative-slope and the yoked LVI conditions.

Method

Subjects

Two male adult humans participated fully in Experiment 2. All of the conditions of their participation were the same as those described above for the subjects in Experiment 1.

Apparatus

The apparatus was the same as that used in Experiment 1.

Procedure

The instructions described above for Experiment 1 were read to each subject prior the first session and were displayed on the computer monitor before the start of each session.

Sessions consisted of three 10-min blocks separated by rest periods that were terminated with a response on the computer keyboard. As in Experiment 1, the subject initiated the start of the session by entering his subject number. Single sessions were conducted on weekdays at approximately the same time of day.

Experimental conditions

Table 2 contains the sequence of conditions and the number of sessions conducted under each for the two subjects. Conditions were changed when mean response rates and within-session patterns of responding were deemed stable via visual inspection of graphical representations of the data. The following are descriptions of each of the 4 experimental conditions.

Variable interval (VI) 30 s and linear variable interval (LVI) 30 s. The conditions were the same as those described above for Experiment 1.

Negative slope (NS) FR n. The contingencies of the negative slope conditions were the same as those described in Experiment 1. As before, the determination was based upon the mean response rates from the final 5 sessions of the LVI condition; the value was selected such that the overall rate of point delivery would equal 0.5 points per min if the response rate of the LVI condition prevailed.

Yoked Linear Variable Interval (Yoked LVI). A linear variable interval schedule was in effect during this condition. The distribution of intervals used to schedule point delivery

Table 2. The sequence of conditions (and the number of sessions conducted under each) for both subjects in Experiment 2 under variable interval (VI), linear variable interval (LVI), negative slope (NS), and yoked LVI schedules of point delivery.

Subject	
S711	S911
VI (24)	VI (29)
LVI (16)	LVI (9)
NS 17 (11)	NS 10 (13)
VI (9)	VI (10)
LVI (4)	VI LH 2s (10)
NS 16 (12)	VI (8)
VI (14)	LVI (7)
LVI (9)	Yoked LVI (13)
Yoked LVI (12)	

was composed of the interpoint intervals obtained under the NS n conditions for each subject, presented in the same sequence in which the points were delivered.

Procedural irregularities

For Subject 711, the response rates during the initial 5-session exposure to the VI condition were very low. In an attempt to increase response rates, relatively shorter IRTs were differentially reinforced manually in Sessions 5 through 10. Also, due to a programming error, the interpoint intervals obtained from Subject 711 under the NS 17 condition were inaccurate, which necessitated an additional negative-slope condition. In an attempt to recover VI response rates before proceeding to the yoked LVI condition, Subject 911 was exposed to a VI 30 s with a 2-s limited hold. The contingencies in effect during this condition were similar to those of the VI 30 s condition except that upon completion of an interval, reinforcement was made available for only 2 s. If a response did not occur within 2 s from the completion of an interval the point delivery was canceled and the next interval began timing.

Results

Molar Feedback Functions

Figure 11 shows the mean rate of point delivery versus the mean response rate from the last 5 sessions under all conditions for Subjects 711 (top panel) and 911 (bottom panel). As in Figure 3 from Experiment 1, these points have been superimposed upon the feedback functions for the respective schedules. Due to the reversal design employed in Experiment 2, the subjects received multiple exposures to the VI and LVI conditions. For clarity, data from similar conditions share the same symbol and the points have been

Figure 11. Mean rate of point delivery versus the mean response rate from the last 5 sessions under all conditions for Subjects 711 (top panel) and 911 (bottom panel). The data have been superimposed upon the molar feedback functions for the respective schedules. Data from similar conditions share the same symbol and the points have been numbered according to the sequence in which were implemented (See Table 2 and the text for details). Note individually scaled axes.

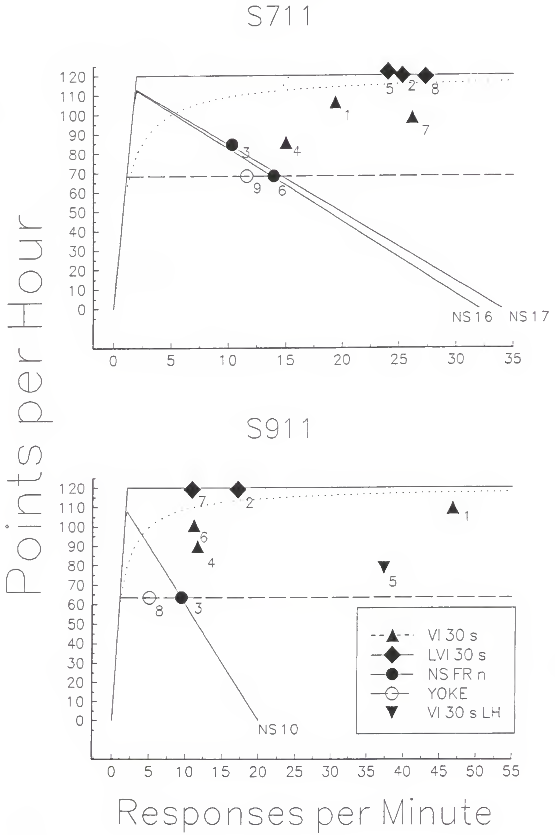
S711
120
110
100
90
80
70
60
50
40
30
20
10
0
0
5
10
15
20
25
30
35
NS 16
NS 17
S911
0
5
10
15
20
25
30
35
40
45
50
55
NS 10
VI 30 s
LVI 30 s
NS FR n
YOKE
VI 30 s LH
Points per Hour
Responses per Minute

numbered according to the sequence in which were implemented (see Table 2). Sensitivity to changes in the contingencies can be assessed by comparing the relative heights and left-right positioning of points across successive conditions. Consider, for example, the relative positions of the points labeled 1 and 2 for Subject 711. These data are from the first exposures to the VI and LVI conditions, respectively. Point 2 lies up and to right of Point 1, indicating that rates of responding and point delivery increased in the LVI condition.

For Subject 711, the steady-state rates of responding and point delivery were consistently lower in the VI conditions (Points 1, 4, & 7) than in the ensuing LVI conditions (Points 2, 5, & 8, respectively). The steady-state rates of responding and point delivery in the NS conditions (Points 3 & 6) were lower than those of the immediately preceding LVI conditions (Points 5 & 8), indicating sensitivity to the negative slope contingencies. The steady-state response rates in the yoked LVI condition (Point 9) were substantially lower than the steady-state rates from the preceding LVI condition (Point 8) and slightly lower than the negative slope condition in which the interpoint intervals were obtained (Point 6).

For Subject 911 (bottom panel), steady-state rates of responding were higher in the first VI condition (Point 1) than in the subsequent LVI condition (Point 2). In the negative-slope condition (Point 3), steady-state rates of responding and point delivery decreased. Upon return to the VI condition (Point 4), steady-state rates of responding increased only slightly. To rectify this, a limited-hold on the availability of points scheduled by the VI 30 s schedule was implemented. Although response rates increased

within this condition (Point 5, See also Figure 12 below), the rates decreased again once the limited hold was discontinued (Point 6). Steady-state response rates decreased slightly in the subsequent LVI condition (Point 7). As with Subject 711, the steady-state response rates in the yoked LVI condition (Point 8) were lower than the steady-state rates from the preceding LVI condition (Point 7) and from the negative slope condition in which the interpoint intervals were obtained (Point 3).

Session-by-Session Patterns

Figure 12 shows the mean response rates (filled circles) and mean rates of point delivery (open circles) across sessions for all conditions for both subjects. Whereas only steady-state data were portrayed in Figure 11, this figure shows transition data and trends in response and reinforcement rates within each condition. For Subject 711 (top panel), response rates were initially very low in the first VI condition. In an attempt to increase response rates, relatively shorter IRTs were differentially reinforced manually in Sessions 5 through 10. Rates of responding and point delivery increased during these sessions. The effect was short lived, however, as response rates decreased throughout the remainder of the condition before reaching a steady state of approximately 20 per minute.

The interpoint intervals used to schedule the point deliveries in the yoked LVI condition were obtained in the NS 16 condition, because the interpoint intervals obtained in the preceding NS 17 condition were inaccurate due to a programming error. Rates of responding and point delivery decreased in this condition relative to those of the immediately preceding LVI condition. In the ensuing VI and LVI conditions, the rates of

Figure 12. The mean response rates (filled circles) and mean rates of point delivery (open circles) across sessions for all conditions for Subjects 711 (top panel) and 911 (bottom panel). Note individually scaled axes.

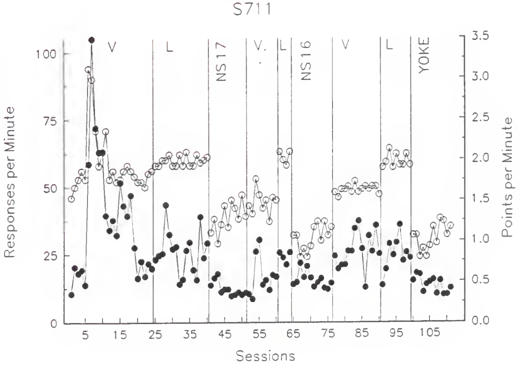
S711
V
L
NS17
V.
L
NS16
V
L
YOKE
Responses per Minute
Points per Minute
Sessions
0
25
50
75
100
0.0
0.5
1.0
1.5
2.0
2.5
3.0
3.5
5
15
25
35
45
55
65
75
85
95
105

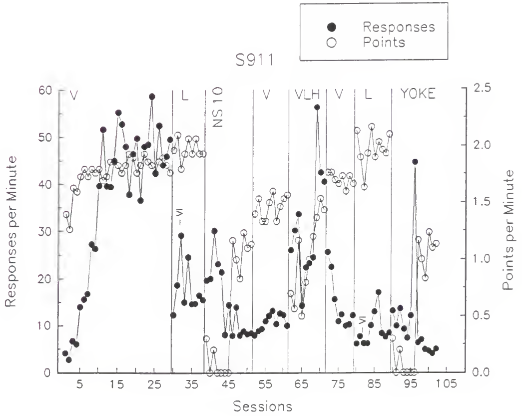
Responses
Points
S911
V
L
NS10
V
VLH
V
L
YOKE
Responses per Minute
Points per Minute
Sessions
0
10
20
30
40
50
60
0.0
0.5
1.0
1.5
2.0
2.5
5
15
25
35
45
55
65
75
85
95
105

responding and point delivery increased to levels comparable to those of the previous exposures, indicating a recovery of baseline performance.

When the yoked LVI condition was implemented, response rates were slightly lower than the rates in the preceding LVI condition, and approximately equal to the rates in the NS 16 condition by the conclusion. The mean response rate from the last 5 sessions under the yoked LVI condition was 42% of the mean response rate from the last 5 sessions of the immediately preceding LVI condition, whereas the mean response rate from the last 5 sessions of the NS 16 condition was 58% of the immediately preceding LVI condition.

Note the high degree of correspondence between the rates of point delivery in the NS 16 and yoked LVI conditions, which is in accord with the programmed contingencies. The only deviations from exact accordance occurred in Sessions 107 and 108. The store was positive at the conclusion of Session 107, with a point carried over to Session 108. As a result, the overall rate of point delivery was slightly lower in Session 107 than in Session 74 (the session to which it was yoked) and slightly higher in Session 108 than in Session 75.

The bottom panel of Figure 12 shows the mean response rates (filled circles) and mean rates of point delivery (open circles) across sessions for all conditions for Subject 911. Early in the initial exposure to the VI condition, response rates were low, but then increased across sessions. Rates of point delivery also increased gradually across the condition. In the subsequent LVI condition, however, response rates decreased considerably while rates of point delivery increased slightly.

In the NS 10 condition, response rates increased across the first three sessions and remained higher than the steady state rates of the preceding LVI condition until Session 43. As a result, no points were delivered in 5 out of the first 7 sessions of the condition. By Session 46, however, response rates had decreased sufficiently to allow the "store" to again become positive. Variability in response rates decreased in the remainder of the condition, with rates of point delivery reaching a steady state of 1.08 per minute.

When the VI conditions were reinstated, the rates of responding and point delivery increased only slightly and were much lower than the rates obtained in the first exposure. In an attempt to recover baseline rates of responding, a 2-s limited hold was implemented. By the conclusion of this condition response rates increased and were comparable to the rates obtained during the initial exposure to the VI condition, but when the VI condition was reinstated, the rates of responding again decreased. Due to time constraints, no additional manipulations to increase response rates were implemented and the experiment progressed, despite the failure to recover baseline response rates completely.

In the first five sessions of the yoked LVI condition, few points were delivered and response rates increased slightly relative to those at the conclusion of the preceding LVI condition. In Session 96, there was a period of high rate responding during the third block, resulting in a considerable increase in session-wide response rate despite an absence of point deliveries. Once the yoked LVI condition had reached the point in the sequence of sessions at which point delivery occurred again, the response rates had decreased. By the conclusion of the condition, response rates were somewhat less than

the rates from the corresponding negative slope sessions. The mean response rate from the last 5 sessions under the yoked LVI condition was 47% of the mean response rate from the last 5 sessions of the immediately preceding LVI condition, whereas the mean response rate from the last 5 sessions of the NS 16 condition was 55% of the immediately preceding LVI condition.

Within-Session Patterns

Sections of representative cumulative records from sessions conducted under NS and yoked LVI conditions are shown in Figure 13 for Subject 711 (top panels) and 911 (bottom panels). Each panel is the segment from the second 10-min block of the session. Under the negative slope conditions, neither subject showed the "pause-run" pattern prevalent for three of the four subjects in Experiment 1. Although relatively long pauses that concluded with point delivery were fairly frequent, they were often separated by periods in which responding occurred at low rates that fluctuated slightly from moment to moment. Under the yoked LVI condition, overall response rates decreased for both subjects, as indicated by the shallower slope of the records on the right. The within-session patterns under the yoked LVI conditions differed for the two subjects, however. For Subject 911, responding occurred at a low, but fairly constant rate throughout the session. For Subject 711, the characteristics of the within-session patterns were similar across the two conditions. Relatively long pauses ending in point delivery became more frequent in the yoked LVI condition, but they were still separated by periods of fluctuating low rate responding. The response patterns of both subjects occasionally produced changes in the temporal distribution of point deliveries.

Figure 13. Portions of representative cumulative response records from sessions conducted under NS and yoked LVI conditions for Subjects 711 and 911. Each panel is the segment of the cumulative record from the second 10 min block of the session.

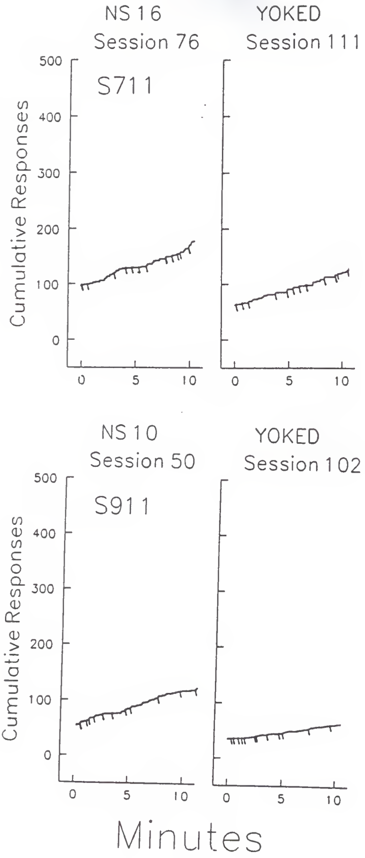
NS 16
Session 76
YOKED
Session 111
S711
Cumulative Responses
500
400
300
200
100
0
0
5
10
NS 10
Session 50
YOKED
Session 102
S911
Minutes

IRTs Preceding Point Delivery

Scatter plots of the IRTs preceding point delivery across sessions for Subjects 711 (top panel) and 911 (bottom panel) are shown in Figure 14. The relevant comparisons in this figure are between the distributions from the yoked LVI conditions and the distributions from the negative slope conditions in which interpoint intervals were obtained (NS 16 and NS 10, for Subjects 711 and 911, respectively). The distributions of IRTs from yoked LVI conditions differ from those of the corresponding negative slope condition for both subjects. For Subject 711, the distribution of IRTs was bimodal in the NS 16 condition. Interestingly, this distribution of IRTs was not readily apparent in the cumulative records for this subject. There was a cluster of IRTs centered around 17 s and another around 0.2 s, and no IRTs were longer than 70 s. In the yoked LVI condition, the distributions maintained a bimodal appearance, but the IRTs were more widely dispersed about the modes. There was a greater frequency of IRTs within the range of 0.5 to 5 s in the yoked LVI condition than in the NS 16 condition and the frequency of relative long IRTs increased, with many exceeding 70 s.

For Subject 911, the distribution of IRTs under the NS 10 condition was not bimodal. At the conclusion of the condition, the IRTs were most densely clustered around 17 s and ranged from approximately 0.25 s to 40 s. In the yoked LVI condition, the frequency of IRTs less than 2 s was considerably less, and the frequency of IRTs exceeding 25 s was higher than in the negative slope condition. Because points were not delivered in some of the early sessions of the NS 10 condition, there were no IRTs preceding point delivery in these sessions or in the corresponding yoked LVI sessions.

Figure 14. Scatter plots of IRTs preceding point delivery across sessions for Subjects 711 (top panel) and 911 (bottom panel). Note individually scaled and logarithmic axes.

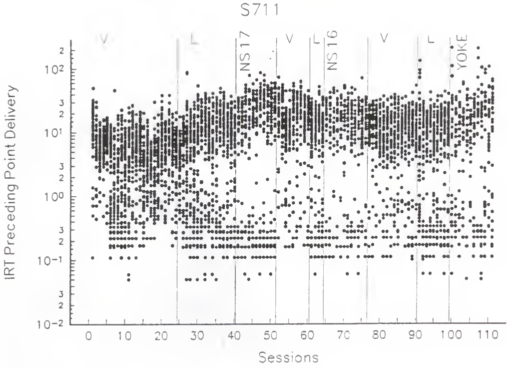
S711
V
L
NS 17
V
L
NS 16
V
L
YOKE
IRT Preceding Point Delivery
Sessions

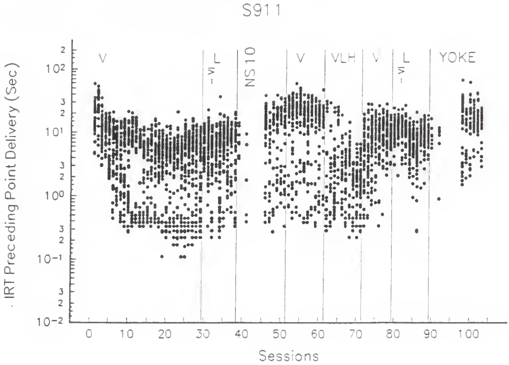
S911
V
L
NS 10
V
VLH
V
L
YOKE
IRT Preceding Point Delivery (Sec)
Sessions

Discussion

The results from the negative slope conditions of Experiment 2 replicate those of Experiment 1. Response rates were lower in the negative slope conditions than in the initial VI baseline conditions, but were not low enough to maximize overall rate of point delivery. Unlike rats (Reed & Schachtman, 1989; 1991) and pigeons (Vaughan, 1982), however, the response rates of both subjects were slightly lower in the yoked LVI condition than in the corresponding negative-slope condition (see Figures 11 & 12). At first glance this might suggest that human behavior is not sensitive to the inverse relationship between rate of responding and rate of point delivery. The results, however, are difficult to interpret because the schedule performances of these two subjects differed from those of the subjects in Experiment 1, suggesting that schedule control may have been partial or incomplete. Specifically, the response rates of these subjects were much lower in the VI and LVI conditions than those of the subjects in the previous experiment. Also, under negative-slope conditions, the distinctive "pause and run" pattern evident in the cumulative records of four of five subjects in Experiment 1 was not apparent for either subject (see Figure 13) and bimodal distributions of IRTs preceding point delivery were obtained for Subject 711 only (see Figure 14). If these characteristics of performance are typical of terminal performance on negative slope schedules, then perhaps the subjects in Experiment 2 did not receive sufficient exposure to the negative slope contingencies.

The reversal design employed in Experiment 2 may also have produced troublesome history effects that complicate interpretation of the results. For Subject 911, response

rates decreased in the first exposure to the LVI condition and, following exposure to the negative slope contingency, it was not possible to recover baseline levels of performance. Exposure to these conditions possibly attenuated sensitivity to additional manipulations of the contingencies. These results may not be too surprising given the prevalence of human behavioral persistence under laboratory conditions (Weiner, 1969, 1970, 1983). As Weiner (1983) has noted, human response patterns tend to persist unless "adverse reinforcement consequences induce change" (p. 521). For Subject 911, the transition from the negative slope condition to the VI condition was met with an increase in rate of reinforcement. These are precisely the circumstances under which one might expect to see persistence of the response patterns established in the earlier condition. When a 2-s limited hold on the availability of points scheduled by the VI was imposed, reinforcement rates decreased initially and then increased as response rate increased. The initial decrease in reinforcement rate was sufficient to induce a change in responding. Rates of point delivery increased when the limited hold was discontinued and remained elevated despite decreases in response rates. Under the yoked LVI condition, previously established patterns of responding were likely to persist because there were no decreases in reinforcement rates to drive changes in response rates.

In a postexperiment questionnaire, Subject 911 reported a contingency description and a performance strategy that were broadly consistent with the contingencies of the negative-slope condition and with his performance. In response to the query;

"Did pressing the red panel influence point delivery?
If so, please describe how as best you can."

Subject 911 wrote;

". . . It seemed to me points were only available if you pressed the red button. But within the tests, it seemed to me that if you pressed the red button too often points were more elusive, which led me to a strategy of restraint in how often I pushed the button."

Although the origin and causal status of this verbal report are dubious, it is consistent with his nonverbal performance, and suggests that at least some aspect of the subject's behavior was sensitive to the negative-slope contingencies. The behavior produced by the negative slope condition (verbal and nonverbal) may have also been occasioned in the yoked LVI condition by the sudden decrease in rate of point delivery. It is then possible that the contingencies of the yoked LVI condition may have sustained these patterns of behavior. That is, Subject 911 may have been sensitive to presence of the subtraction contingency in the negative slope condition, but not sensitive to the absence of the contingency in the yoked LVI condition. Performance in the yoked LVI condition may have been an instance of what Shimoff, Catania, and Mathews (1986) referred to as "pseudosensitivity" to the contingencies--rule-governed response patterns in accord with the contingencies. In the present case, the response rates on the yoked LVI may have been produced directly by the rate of point delivery. Alternatively, they may have been influenced by verbal descriptions of the contingencies generated under the negative slope conditions and occasioned in the yoked LVI condition.

In sum, the results of Experiment 2 are inconclusive with respect to the question of whether human behavior is sensitive to the presence or absence of the subtraction contingency. At least some of the problems may have been circumvented if the interpoint intervals were yoked across components of a multiple schedule rather than

across experimental conditions. The addition of discriminative stimuli for each of the different conditions may have helped foster sensitivity to the differences between the yoked and the negative slope conditions, and minimized the persistence of preestablished response patterns.

GENERAL DISCUSSION

The overall pattern of results from both experiments suggests that human behavior is sensitive to the contingencies of negative-slope schedules. This was evidenced primarily by the decreases in overall rates of responding across negative-slope conditions in Experiment 1 (See Figure 3). In 22 of 23 negative-slope conditions across subjects, the steady-state response rates were lower than those of the preceding condition where the subtraction contingency was less stringent, indicating sensitivity to changes in the contingencies. These findings were partially replicated in Experiment 2 in that overall response rates in the negative slope conditions were lower than those of the initial VI conditions. The development of similar within-session response patterns across several subjects also indicates sensitivity to the negative-slope contingencies.

Implications for Maximization Accounts

Sensitivity to the inverse relationship between overall response rate and overall reinforcement rate was limited. Complete sensitivity to this relationship would have been evidenced by overall response rates that were consistently around 2 responses per minute. At this rate, all available points would have been collected in a timely fashion and the impact of the subtraction ratio on overall rate of point delivery would have been minimized. The results of Experiments 1 and 2 are therefore not in accord with

interpretations based strictly upon reinforcement maximization. In Figures 3 and 11, the data from the negative slope conditions do not lie near the apexes of the molar feedback functions for the respective schedules. By failing to respond at or around 2 responses per minute, the subjects failed to maximize overall rate of point delivery or to minimize the number of responses per point.

The human subjects in the present experiments fared better with respect to overall reinforcement density maximization than the pigeons in Vaughan and Miller's (1984) experiment, however. Across subjects, the response rates tended to fall closer to the apexes of the feedback function than those of the pigeons. Moreover, reinforcement rates occasionally increased for humans as the subtraction ratios became more stringent, whereas the pigeons tended to perform poorly at the lower subtraction ratios.

These results complement the results of experiments examining performance on concurrent response-based and time-based schedules (Silberberg, Thomas, & Berendzen, 1991; Savastano & Fantino; 1994). On these procedures, reinforcement rate is maximized by responding mostly on the response-based alternative; the time-based schedule should be sampled only briefly and occasionally in order to collect reinforcers as they are set up. Both pigeons (Heyman & Herrnstein, 1986) and humans (Silberberg et al., 1991; Savastano & Fantino, 1994) allocate too much time to the time-based alternative and overall reinforcement rate is not maximized. Humans, however, tend to show a greater bias for the ratio alternative than do pigeons.

The present results are at odds with those of prior experiments examining "self-control" (Logue et al., 1986) and choice in situations of diminishing returns (Hackenberg

& Axtell, 1992; Jacobs & Hackenberg, 1996; Wanchisen, et al., 1992) in humans, which are generally consistent with accounts based upon the maximization of overall reinforcement density. Indeed, deviations from maximization have become the exception rather than the rule in laboratory experiments with adult human subjects. It may be desirable to search these few exceptional cases for common features. This task is complicated, however, by significant methodological differences across experiments. In an experiment by Wasserman and Neunaber (1986), for example, pressing a telegraph key was maintained in student volunteers by contiguous point presentations even when pressing also decreased the overall rate of point delivery. These results suggest that human behavior can under certain circumstances be maintained by short-term consequences at the expense of long-term outcomes. It is important to note, however, that exposure to the contingencies was brief (1 session) and the points were not exchangeable for other reinforcers (i.e., they were merely "instructed" reinforcers). The effects can therefore be questioned on the grounds that the subjects did not have sufficient exposure to the contingencies to contact the long-term consequences, and that the point presentations were weak reinforcers (if indeed they were reinforcers at all).

Implications for Matching Accounts

Matching-based interpretations do not account well for the steady-state data either. As discussed above, the steady-state data from Experiment 1 portrayed in Figure 4 do not conform to Equation 2. This is in contrast with the results of several other studies examining human performance on VI schedules of reinforcement presented either alone or concurrently (Baum, 1975; Bradshaw, Szabadi, & Bevan, 1976, 1977, 1979;

McDowell & Wood, 1984; Shroeder & Holland, 1969: but see, Horne & Lowe, 1993; Schmitt, 1974).

In two studies by Bradshaw and colleagues (1977, 1979) the effects of punishment on human variable-interval performance was examined. These studies warrant mention because, as in the present study, an inverse relationship between overall response rate and overall reinforcement rate was established. In these studies, button pressing was maintained in humans by points exchangeable for money. Point deliveries were arranged by a multiple schedule comprised of five different VI component schedules. Point delivery was signaled by a 100 ms illumination of a green light and the addition of one point to a counter. Under some conditions, a variable-ratio 34 schedule of point loss was in effect conjointly with the VI schedules. Every 34th response on average produced point loss signaled by a 100 ms illumination of a red lamp and the subtraction of one point from the counter.

Although there are many differences between this procedure and that of the negative slope conditions of the present study, both established an inverse relationship between overall response rate and overall reinforcement rate. The net result was the same: higher rates of responding resulted in lower rates of reinforcement overall. The major difference between the two procedures lies in the discriminability of point loss. In the Bradshaw et al. studies (1977, 1979), point loss was signaled, whereas the point omissions in the present study were unsignaled.

As in the negative-slope conditions of the present study, response rates in the punishment conditions of the Bradshaw et al. studies (1977, 1979) were lower than

baseline rates. Unlike the results of Experiment 1, however, response rates in the punishment conditions varied as a function of reinforcement rate in a manner consistent with Equation 1. Relative to the data obtained in the unpunished conditions, the presence of the punishment contingency decreased $\underline{K}$, the maximum rate of responding, and increased $\underline{r}_o$, the rate of reinforcement from unprogrammed sources. It would have been interesting to see if the VR schedule of point loss controlled within-session patterns of responding in a manner similar to the subtraction contingency of the negative slope conditions of the present experiment. Unfortunately, cumulative response records were not published.

Although the results of these and other studies by Bradshaw and colleagues lend support to the applicability of matching-based interpretations to human behavior, one aspect of the procedure complicates interpretation of the findings. As Horne and Lowe (1993) pointed out, the discriminative stimuli that accompanied the different component schedules may have been functioning as implicit instructions, indicating differences in the schedule values. The stimuli consisted of a series of lights across the top of the apparatus. The left to right positioning of the lamps covaried with the VI schedule values; the light correlated with the richest schedule was on the extreme right, that correlated with the leanest schedule on the extreme left, with a graded series of schedule values in between.

Horne and Lowe (1993) performed a systematic replication of the Bradshaw et al. (1976) study and produced evidence suggesting that the results could be explained on the basis of implicit ordinal stimuli. In a series of experiments, they found that matching was

frequently obtained when implicit or explicit ordinal cues were used as discriminative stimuli, but significant deviations from matching occurred when different geometrical shapes were projected onto a response panel as discriminative stimuli. Such results suggest that the performances in the Bradshaw et al. studies may have been verbally mediated through implicit instructions.

The Role of Verbal Mediation

Regarding verbal mediation, a possible explanation for the disparity between the results of the present experiment and those of the "self-control" procedures is that responding under the present circumstances was less influenced by verbal behavior. This research was undertaken, in part, to examine human behavior in an environment wherein the relationship between responding and point delivery was difficult to describe. It was hypothesized that the use of subtle contingencies of reinforcement would minimize the influence of verbal behavior on nonverbal performance. Although one subject (Subject 911, Exp. 2) offered a general description the relationship between rate of point delivery and rate of responding, none of the subjects were able to specify a maximization strategy as the subjects in the "self-control" experiments could (Logue et al., 1986). The aperiodicity of point delivery arranged by the present procedures may have been sufficiently subtle to preclude accurate descriptions, thereby eluding the influence of verbal mediation to a significant degree.

Because the present study was not designed to control verbal behavior nor to measure its relation to in nonverbal performance, it is not possible to rule out verbal mediation. To do so would require recording ongoing verbal behavior and its relation to the

nonverbal performance it accompanies. To avoid creating descriptions of the contingencies that might not otherwise have occurred, however, verbal reports were limited to a postexperimental questionnaire. Given the preliminary nature of the present research, it seemed prudent to examine contingency sensitivity with minimal verbal involvement before approaching the more complex interactions between verbal and nonverbal performance. A more precise characterization of the relationship between verbal behavior and performance under negative-slope schedules awaits further research.

The Structure of Negative-Slope Schedule Performance

Although the present results are not well accounted for by interpretations based upon either the matching law or maximization, it is possible that order exists at a more local time scale. Reed and Schachtman (1991), for example, suggested that the difference between response rates on negative-slope schedules and yoked LVI schedules may be due to the differential reinforcement of relatively longer IRTs under the negative slope contingencies. In their experiment, interreinforcement intervals were yoked across negative slope and LVI components of a multiple schedule. They also required that the IRT preceding reinforcement in the yoked component fall within 1 s of that obtained in the negative slope component. Overall response rates were more similar across the two components when the IRT requirement was present than when it was absent. While this result lends indirect support to the hypothesis that negative slope schedules differentially reinforce longer IRTs, the results are far from unequivocal.

More convincing evidence would have been provided by examining differences in the IRT distributions in the two conditions. Unfortunately, none of the prior studies of

negative slope schedule performance provide data of this sort. In the present study, the IRTs preceding point deliveries were collected across all conditions for 6 of 7 subjects. The distribution of IRTs preceding point delivery became bimodal and the range increased under the negative slope conditions (See Figures 8, 9, & 14) in 4 of these 6 subjects. These data are consistent with Reed and Schachtman's (1991) hypothesis that relatively longer IRTs are differentially reinforced under negative slope schedules.

The bimodal distributions obtained in the present study reveal something that Reed and Schachtman's analysis could not. Merely stating that longer IRTs are reinforced under the negative slope schedules does not address the manner in which the distribution of IRTs changes. Such a statement would be as applicable to a shift in a unimodal distribution as it is to the present case. The fact that bimodal distributions of IRTs preceding point delivery occurred for several subjects suggests something about the structure of negative-slope schedule performance in general. Such performance appears to be characterized by alternating periods of pausing and high-rate responding. Verifying this relationship quantitatively would have required the collection of every IRT and the examination of shifts in the distributions across conditions. Unfortunately, this was not feasible in the present study. Short of recording every IRT, one can still gain some important insights regarding the moment-to-moment characteristics of negative slope schedule performance by examining cumulative response records.

Although cumulative records are not published as frequently as they once were (Skinner, 1976), they are still useful for assessing schedule control by portraying schedule-typical moment-to-moment fluctuations in responding (Ferster & Skinner,

1957; Zeiler, 1984). In all studies employing negative-slope schedules to date (Ettinger, et al., 1987; Reed & Schachtman, 1989; 1991; Vaughan & Miller, 1984) however, there has been only one cumulative record published (Ettinger, et al., 1987) and it is not particularly useful. In that study, lever pressing was maintained in rats by an LVI 90 s schedule of food presentation, with food deliveries conjointly omitted according to an FR 20 schedule. A composite cumulative record of the first 20 responses following food presentation averaged across the entire session was published. As a result of the limited range of data presented and the curve smoothing effects of the transform, the record is of little use in discerning the characteristic moment-to-moment fluctuations in responding generated by the schedule. The absence of cumulative records of negative-slope performance in the literature makes it difficult to determine if adequate schedule control had been achieved.

In the present study, cumulative records were obtained for all subjects under all conditions. Across the two experiments, the negative slope schedules generated a "pause and run" pattern of responding in 4 of 7 subjects. The onset of responding was usually met with a series of point deliveries until the "store" was depleted. Responding often continued after the "store" was depleted and additional points were occasionally delivered. This pattern differs considerably from that typical of VI performance, characterized by responding at a constant rate with little pausing. For the four subjects (Subjects 111, 888, 211, and 521) showing this pattern, typical VI patterning was established in the initial VI condition, was maintained throughout the LVI condition, then transitioned to the "pause and run" pattern across negative-slope conditions. The

development of this pattern provides further evidence that the subjects' behavior was sensitive to the subtraction contingency. If responding was simply tracking changes in overall rate of point delivery, there is no reason to suppose that response patterns would differ across the VI and negative-slope conditions. Indeed, one might expect differences to be seen primarily in the overall rates, not in the local patterning.

Of the 3 subjects (Subjects 999, 711, and 911) that did not show the "pause-run" pattern robustly, all responded at much lower rates in the baseline VI and LVI conditions than the other subjects, with two (Subjects 999 and 911) showing sensitivity to the VI-LVI difference. Moreover, these subjects generally received less exposure overall to the negative slope schedules than some of the other subjects. It is possible that these subjects did not receive sufficient exposure to the negative slope schedules for the "pause-run" pattern to develop.

It is also possible that the "pause-run" pattern of responding depends on the baseline response rate. The response patterns of humans on fixed interval schedules, for example, generally fall into one of two categories: low rate, characterized by long pauses with a few responses at the end of the interval, or high rate, characterized by infrequent pausing and virtual insensitivity to schedule changes (Laties & Weiss, 1963; Lippman & Meyer, 1967). Weiner (1969) has demonstrated that a prior experimental history can influence the pattern of FI responding. Subjects initially exposed to an IRT > t s schedule tended to respond at a low rate when subsequently exposed to a series of FI schedules, whereas those initially exposed to a fixed-ratio schedule tended to respond at a high constant rate under the FI schedules. It is possible that the subjects' preexperimental histories and/or

their low-rate histories on the baseline VI and LVI conditions modulated sensitivity to the negative-slope conditions of the present experiment.

Cumulative response records may be a useful tool for the identification of response units created and maintained under negative slope schedules. The "pause and run" response pattern may be one type of unit characteristic of negative-slope schedule performance. Although the reliability of this pattern across procedures and species needs to be established through replication, examination of these regularities in responding may prove useful in the identification of appropriate behavioral units and the development of quantitative descriptions that improve prediction and control. In future research, every IRT should be recorded to examine cyclical fluctuations in response rate as they are related to cyclical variations in the temporal distribution of reinforcement. By collecting every IRT, behavioral units across different time scales can be examined. The "pause and run" pattern itself, for example, may be the aggregate effect of differential reinforcement acting on different classes of IRTs, or on some local property of responding. Mathematical models formulated to account for the local dynamics of behavior (e.g., Killeen, 1994) may one day be brought to bear on negative slope performance as well.

SUMMARY AND CONCLUSIONS

Previous research with human subjects has implicated verbal behavior as a possible determinant of between-species differences in sensitivity to temporally remote outcomes. The results of the present study tentatively support this position and suggest one possible method for minimizing the influence of verbal behavior on nonverbal performance. Human behavior was examined under contingencies that did not give rise to accurate, quantitative descriptions of the relationship between responding and point delivery and results more similar to those of nonhumans were obtained. The results suggest that human behavior may be more likely to be controlled by immediate outcomes when it is difficult to describe the relationship between responding and consequences. Due to limitations in the resolution of data collection, however, this study should be regarded as preliminary. Because assessment of verbal behavior was limited in the present study, additional research is needed to evaluate the role of verbal behavior in establishing and maintaining sensitivity to temporally remote outcomes. Additional research is also necessary to provide a precise quantitative account of negative-slope schedule performance.

REFERENCES

Ainslie, G. W. (1974). Impulse control in pigeons. Journal of the Experimental Analysis of Behavior, 21, 485-489.

Baum, W. M. (1973). The correlation-based law of effect . Journal of the Experimental Analysis of Behavior, 20, 137-153.

Baum, W. M. (1975). Time allocation in human vigilance. Journal of the Experimental Analysis of Behavior, 23, 45-53.

Bradshaw, C. M., Szabadi, E, & Bevan, P. (1976). Behavior of humans in variable-interval schedules of reinforcement. Journal of the Experimental Analysis of Behavior, 27, 275-279.

Bradshaw, C. M., Szabadi, E, & Bevan, P. (1977). Effect of punishment on human variable-interval performance. Journal of the Experimental Analysis of Behavior, 27, 275-279.

Bradshaw, C. M., Szabadi, E, & Bevan, P. (1979). The effect of punishment on free-operant choice in humans. Journal of the Experimental Analysis of Behavior, 31, 71-81.

Catania, A. C., Mathews, B. A., & Shimoff, E. (1982). Instructed versus shaped human verbal behavior: Interactions with nonverbal responding. Journal of the Experimental Analysis of Behavior, 38, 233-248.

Catania, A. C., & Reynolds, G. S. (1968). A quantitative analysis of the responding maintained by interval schedules of reinforcement. Journal of the Experimental Analysis of Behavior, 11, 327-383.

Cerutti, D. T. (1989). Discriminative theory of rule-governed behavior. Journal of the Experimental Analysis of Behavior, 51, 259-276.

Ettinger, R. H., Ried, A. K., & Staddon, J. E. R. (1987). Sensitivity to Molar Feedback Functions: A Test of Molar Optimality Theory. Journal of Experimental Psychology: Animal Behavior Processes, 13, 366-375.

Ferster, C. B., & Skinner, B.F. (1957). Schedules of reinforcement. New York: Appleton-Century-Crofts.

Fleshler M., & Hoffman, H. S., (1962). A progression for generating variable-interval schedules. Journal of the Experimental Analysis of Behavior, 5, 529-530.

Hackenberg, T. D., & Axtell, S. A. M. (1993). Humans' choices in situations of time-based diminishing returns. Journal of the Experimental Analysis of Behavior, 59, 445-470.

Hackenberg, T. D., & Hineline, P. N. (1992). Choice in situations of diminishing returns: Immediate versus delayed consequences of action. Journal of the Experimental Analysis of Behavior, 57, 67-80.

Herrnstein, R. J. (1970). On the law of effect. Journal of the Experimental Analysis of Behavior, 13, 243-266.

Heyman, G. M., & Herrnstein, R. J., (1986). More concurrent interval-ratio schedules: A replication and review. Journal of the Experimental Analysis of Behavior, 46, 331-351.

Hineline, P. N., & Sodetz, F. J. (1987). Appetitive and aversive schedule preferences: Schedule transitions as intervening events. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), Quantitative analyses of behavior: Vol. 5. The effect of delay and intervening events on reinforcement value (pp. 141-157). Hillsdale, NJ: Erlbaum.

Hodos, W., & Trumbule, G. H. (1967). Strategies of schedule preference in chimpanzees. Journal of the Experimental Analysis of Behavior, 10, 503-514.

Horne, P. J., & Lowe, C. F. (1993). Determinants of human performance on concurrent schedules. Journal of the Experimental Analysis of Behavior, 59, 29-60.

Horne, P. J., & Lowe, C. F. (1996). On the origins of naming and other symbolic behavior. Journal of the Experimental Analysis of Behavior, 65, 185-241.

Hyten, C., Madden, G. J., and Field, D. P. (1994). Exchange delays and impulsive choice in adult humans. Journal of the Experimental Analysis of Behavior, 62, 225-233.

Jackson, K., & Hackenberg, T. D. (1996). Token reinforcement, choice, and self-control in pigeons. Journal of the Experimental Analysis of Behavior, 66, 29-49.

Jacobs, E. A., & Hackenberg, T. D., (1996). Humans' choices in situations of diminishing returns: Effects of fixed-interval duration and progressive-interval step size. Journal of the Experimental Analysis of Behavior, 65, 5-19.

Logue, A. W. (1988). Research on self-control: An integrating framework. Behavioral & Brain Sciences, 11, 665-709.

Logue, A.W., & King, G.R. (1991). Self-control and impulsiveness in adult humans when food is the reinforcer. Appetite, 17, 105-120.

Logue, A.W., Pena-Correal, T.E., Rodriguez, M.L., & Kabela, E. (1986). Self-control in adult humans: Variation in positive reinforcer amount and delay. Journal of the Experimental Analysis of Behavior, 46, 159-173.

Killeen P. (1994). Mathematical principles of reinforcement. Behavioral and Brain Science, 17, 105-172.

Mazur, J. E., & Vaughan, W. Jr. (1987). Molar optimization versus delayed reinforcement as explanations of choice between fixed-ratio and progressive-ratio schedules. Journal of the Experimental Analysis of Behavior, 48, 251-261.

McDowell, J.J., & Wood, H.M. (1984). Confirmation of linear systems theory prediction: Changes in Herrnstein's k as a function of changes in reinforcer magnitude. Journal of the Experimental Analysis of Behavior, 41, 183-192.

Navarick, D. J. (1982). Negative reinforcement and choice in humans. Learning and Motivation, 13, 361-377.

Prelec, D. (1982). Matching, maximizing, and the hyperbolic reinforcement feedback function. Psychological Review, 89, 189 - 230.

Rachlin, H., & Green, L. (1972). Commitment, choice, and self-control. Journal of the Experimental Analysis of Behavior, 17, 15-22.

Reed, P., & Schachtman, T. R. (1991). Instrumental performance on negative schedules. The Quarterly Journal of Experimental Psychology, 43b, 177-197.

Reed, P., & Schachtman, T. R. (1989). Instrumental responding by rats on free operant contingencies with components that schedule response-dependent reinforcer omission. Animal Learning and Behavior, 17, 328-338.

Rosenfarb, I. S., Newland, M. C., Brannon, S. E., & Howey, D. S. (1992). Effects of self-generated rules on the development of schedule-controlled behavior. Journal of the Experimental Analysis of Behavior, 58, 107-121.

Savastano, H., & Fantino, E. F. (1994). Human choice in concurrent ratio-interval schedules of reinforcement. Journal of the Experimental Analysis of Behavior, 61, 453-463.

Schmitt, D.R. (1974). Effects of reinforcement rate and reinforcer magnitude on choice behavior of humans. Journal of the Experimental Analysis of Behavior, 21, 409-419.

Schroeder, S.R., & Holland, J.G., (1969). Reinforcement of eye movement with concurrent schedules. Journal of the Experimental Analysis of Behavior, 12, 897-903.

Shimoff, E., Matthews, B.A., & Catania, A.C. (1986). Human operant performance: Sensitivity and pseudosensitivity to contingencies. Journal of the Experimental Analysis of Behavior, 46, 149-157.

Silberberg, A., Thomas, J. R., & Berendzen, N. (1991). Human choice on concurrent variable-interval, variable-ratio schedules. Journal of the Experimental Analysis of Behavior, 56, 575-584.

Skinner, B. F. (1957). Verbal behavior, New York: Appleton-Century-Crofts.

Skinner, B. F. (1969). An operant analysis of problem solving. In K. MacCorquodale, G. Lindzey, & K. E. Clark (Eds.), Cumulative record (pp. 295-313).

Skinner, B. F. (1976). Farewell, my lovely. Journal of the Experimental Analysis of Behavior, 25, 218.

Solnick, J. V., Kannenberg, C.H., Eckerman, D. A., & Waller, M. B. (1980). An experimental analysis of impulse control in humans, Learning and Motivation, 11, 61-77.

Sonuga-Barke, E.J.S., Lea, S.E.G., & Webley, P.(1989). The development of adaptive choice in a self-control paradigm. Journal of the Experimental Analysis of Behavior, 51, 77-85.

Vaughan, W., Jr. (1982). Choice and the Rescorla-Wagner model. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.) Quantitative analyses of behavior: Vol. 2. Matching and maximizing accounts (pp. 263 -279). Cambridge, MA: Ballinger.

Vaughan, W., Jr. & Miller, H. L. (1984). Optimization versus response strength accounts of behavior. Journal of the Experimental Analysis of Behavior, 42, 337-348.

Wanchisen, B. A., Tatham, T. A., & Hineline, P. N. (1988). Pigeons' choices in situations of diminishing returns: Fixed-versus progressive-ratio schedules. Journal of the Experimental Analysis of Behavior, 50, 375-394.

Wanchisen, B. A., Tatham, T. A., & Hineline, P. N. (1992). Human choice in "counterintuitive" situations: Fixed- versus progressive schedules. Journal of the Experimental Analysis of Behavior, 58, 67-85.

Wasserman, E. A., & Neunaber, D. J., (1986). College students' responding to and rating of contingency relations: The role of temporal contiguity. Journal of the Experimental Analysis of Behavior, 46, 15-35.

Weiner, H. (1969). Controlling human fixed-interval performance. Journal of the Experimental Analysis of Behavior, 12, 349-373.

Weiner, H. (1970). Human behavioral persistence. The Psychological Record, 20, 445-456.

Weiner, H. (1983). Some thoughts on discrepant human-animal performances under schedules of reinforcement. The Psychological Record, 33, 521-532.

Zeiler, M.D. (1984). Schedules of reinforcement: The sleeping giant. Journal of the Experimental Analysis of Behavior, 42, 485-493.

BIOGRAPHICAL SKETCH

Eric A. Jacobs was born in South Amboy, New Jersey, on May 17, 1967. When he was 12 years old, the Jacobs family moved to Florida. After graduating from Palm Beach Gardens Community High School in 1985, Eric attended the University of Florida where he studied psychology. In 1990, he received a Bachelor of Science degree and began graduate studies under the direction of Timothy D. Hackenberg, Ph.D. In 1994, he was awarded the degree of Master of Science. Eric currently resides with his wife, Angela Rabb Jacobs, in Burlington, Vermont, where he serves as an Assistant Project Director for the Substance Abuse Treatment Center at the University of Vermont. Eric longs for the camaraderie of colleagues at the University of Florida, for the simplicity of microswitches, for the clatter of cumulative response recorders, and for the elegance of within-subject designs.

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Timothy D. Hackenberg, Chair
Associate Professor of Psychology

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Marc N. Branch
Professor of Psychology

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Brian A. Iwata
Professor of Psychology

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Henry S. Pennypacker
Professor of Psychology

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Donald J. Stehouwer
Professor of Psychology

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Mark H. Lewis
Professor of Neuroscience

This dissertation was submitted to the Graduate Faculty of the Department of Psychology in the College of Liberal Arts and Sciences and to the Graduate School and was accepted as partial fulfillment of the requirements for the degree of Doctor of Philosophy.

August, 1997

Dean, Graduate School